Multidimensional Manhattan Sampling and Reconstruction

Matthew A. Prelee and David L. Neuhoff

Abstract

This paper introduces Manhattan sampling in two and higher dimensions, and proves sampling theorems. In two dimensions, Manhattan sampling, which takes samples densely along a Manhattan grid of lines, can be viewed as sampling on the union of two rectangular lattices, one dense horizontally, the other vertically, with the coarse spacing of each being a multiple of the fine spacing of the other. The sampling theorem shows that images bandlimited to the union of the Nyquist regions of the two rectangular lattices can be recovered from their Manhattan samples, and an efficient procedure for doing so is given. Such recovery is possible even though there is overlap among the spectral replicas induced by Manhattan sampling.

In three and higher dimensions, there are many possible configurations for Manhattan sampling, each consisting of the union of special rectangular lattices called bi-step lattices. This paper identifies them, proves a sampling theorem showing that images bandlimited to the union of the Nyquist regions of the bi-step rectangular lattices are recoverable from Manhattan samples, and presents an efficient onion-peeling procedure for doing so. Furthermore, it develops a special representation for the bi-step lattices and an algebra with nice properties. It is also shown that the set of reconstructable images is maximal in the Landau sense.

While most of the paper deals with continuous-space images, Manhattan sampling of discrete-space images is also considered, for infinite, as well as finite, support images.

Index Terms

image sampling, lattice sampling, Landau sampling rate, nonuniform periodic sampling.

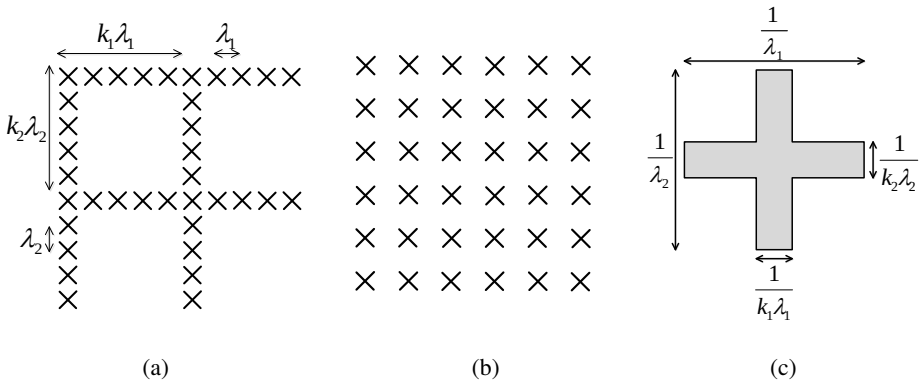

Fig. 1. (a) 2D Manhattan-grid sampling sites with parameters $k_1 = k_2 = 5$ and $\lambda_1 = \lambda_2$. (b) Square lattice sampling at the same density. (c) Cross-shaped frequency support (centered at the origin) of images recoverable with Manhattan sampling.

I. Introduction

In the two-dimensional (2D) setting, *Manhattan sampling* (or *M-sampling* for short) is a recently proposed form of image sampling in which data is taken along evenly spaced rows and columns; the set of sample locations will be called a *Manhattan grid*. In particular, as illustrated in Fig. 1(a), given sampling intervals $\lambda_1, \lambda_2 > 0$ and integers $k_1, k_2 > 1$, samples are taken at intervals of λ_1 along horizontal rows spaced $k_2\lambda_2$ apart, and also at intervals of λ_2 along vertical columns spaced $k_1\lambda_1$ apart.

Manhattan sampling has been used to good effect in both lossy and lossless bilevel image compression [1]–[3]. These methods losslessly compress the samples in a Manhattan grid, for example with arithmetic coding (AC), which can be done with very few bits per M-sample because the samples are closely spaced and, hence, highly correlated. For lossless compression, the other pixels are then AC encoded, conditioned on those in the Manhattan grid, while for lossy compression there is no further encoding, and the decoder estimates the remaining pixels from those in the Manhattan grid. Markov random field models have been used to guide both the arithmetic coding and the estimation.

M-sampling has also been proposed [4], [5], [7], [8] as a new approach to sampling grayscale images and other two-dimensional fields, with the motivations that (a) dense sampling along lines might capture edge transitions more completely than conventional lattice sampling with the same density, (b) sensor networks with a Manhattan deployment geometry need less power or less wire to transmit data than

The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 (email: mprelee@umich.edu; neuhoff@eecs.umich.edu).

This work was supported by NSF Grant CCF 0830438. Portions were presented at IEEE ICASSP 2012 and ITA 2014.

conventional lattice or random deployments at the same density [7], [8], and (c) there are physical scenarios for which M-sampling is far more natural than traditional lattice sampling, such as when sampling from a moving vehicle, e.g., a ship sampling oxygen levels in a body of water. Similarly motivated by sampling from vehicles, the recent related work of Unnikrishnan and Vetterli [9], [10] considers sampling continuously along a grid of lines, i.e., with asymptotically large sampling rate.

Methods for approximately reconstructing typical (non-bandlimited) images from M-samples have been developed in [4]–[6]. The present paper focuses on identifying a bandlimited set of images that can be perfectly reconstructed, as well as efficient methods for doing so.

Manhattan sampling with parameters $\lambda_1, \lambda_2, k_1, k_2$ can be viewed as sampling on the union of the horizontally dense rectangular lattice consisting of all locations of the form $(n_1\lambda_1, n_2k_2\lambda_2)$, where n_1, n_2 are arbitrary integers, and the similarly defined vertically dense rectangular lattice consisting of all locations of the form $(n_1k_1\lambda_1, n_2\lambda_2)$. For brevity, we call these the *horizontal* and *vertical* lattices, respectively.

By the conventional 2D sampling theorem [11] (see also [12, p. 72], [13, Chap. 3], [14, p. 43]), the samples on the horizontal lattice are sufficient to distinguish and reconstruct any image bandlimited to the Nyquist region $\left\{(u,v) : |u| < \frac{1}{2\lambda_1}, |v| < \frac{1}{2k_2\lambda_2}\right\}$. Likewise the samples on the vertical lattice are sufficient to distinguish and reconstruct any image bandlimited to the Nyquist region $\left\{(u,v) : |u| < \frac{1}{2k_1\lambda_1}, |v| < \frac{1}{2\lambda_2}\right\}$. Each of these samplings is maximally efficient in the Landau sense [15] that their sampling densities are as small as the area of the Nyquist region. Equivalently, the set of images bandlimited to the Nyquist region is maximal for the given sampling scheme.

The first result of the present paper is a sampling theorem in Section III showing that images bandlimited to the union of these two Nyquist regions can be reconstructed from their samples on the union of the two rectangular lattices, i.e., on the Manhattan grid, and an efficient procedure for doing so is given. It is also shown that the images bandlimited in this way form a maximal reconstructable set for the Manhattan grid samples. As illustrated in Fig. 1(c), the union of the two Nyquist regions is the cross-shaped *Manhattan region*. We say that images whose spectra are confined to such a region are *Manhattan-bandlimited*. Given the relevance of Manhattan-bandlimiting, a figure in Section III will display the effect of several instances of such on a typical image.

The principal goals of the remainder of the paper are to formulate M-sampling in three and higher dimensions, and to derive a sampling theorem and a reconstruction procedure. M-sampling in three dimensions can be motivated by the need to spatially sample a three-dimensional volume with a vehicle, or to spatio-temporally sample a two-dimensional region, as in video, or a spatio-temporal sensor network.

Four-dimensional sampling can be motivated by the need for spatio-temporal sampling of a three-dimensional spatial region.

In three and higher dimensions, M-sampling can take a variety of forms. In order to describe two of these in three dimensions, consider the partition of 3D space into $k_1\lambda_1 \times k_2\lambda_2 \times k_3\lambda_3$ orthotopes (3D rectangles). As illustrated in Fig. 4(a), one form of M-sampling takes samples uniformly along each edge of each of these orthotopes — with spacing λ_i along edges parallel to axis i. Another form (Fig. 4(c)) takes samples uniformly on each face of each orthotope — with the samples on the face orthogonal to axis i taken according to a $\lambda_j \times \lambda_k$ rectangular lattice, where j and k denote the other dimensions. In other words, the first form samples densely along lines and the second samples densely along hyperplanes. Neither of these takes samples in the interior of any of the aforementioned orthotopes.

More generally, as described in Section IV, M-sampling in an arbitrary dimension d is defined as taking samples on the union of some collection of d-dimensional *bi-step* lattices, which are rectangular lattices defined by step sizes that in dimension i are restricted to λ_i or $k_i\lambda_i$. Thus, there are many possible M-samplings in d dimensions, even when λ_i's and k_i's are fixed. We call such unions of d-dimensional bi-step lattices *Manhattan sets*.

The main results of Sec. IV are (a) a sampling theorem showing that images bandlimited to the union of the Nyquist regions of the d-dimensional bi-step lattices comprising the Manhattan set can be distinguished by their M-samples, (b) efficient, onion-peeling procedures for perfectly reconstructing d-dimensional images, bandlimited as in (a), from their M-samples (one in frequency domain and one in spatial domain), and (c) a proof that the set of such bandlimited images is maximal in the Landau sense.

The development of the sampling theorem and reconstruction procedures are enabled by an efficient parametrization of a bi-step lattice (with a given set of λ_i's and k_i's) by a binary vector $\mathbf{b} = (b_1, \ldots, b_d)$ indicating the dimensions i along which the spacing between lattice points is the smaller value, λ_i, rather than the larger value, $k_i\lambda_i$. This enables any Manhattan set to be compactly described by a finite set of $\mathbf{b}_i$'s (in addition to the λ_i's and k_i's). A number of properties and relationships are enabled by this parametrization. For example, the computation of the density of a d-dimensional Manhattan set is enabled by a spatial partition whose 2^d atoms are indexed by $\mathbf{b}$'s. Similarly, the onion-peeling reconstruction procedures mentioned previously are keyed to a partition of frequency space whose 2^d atoms are indexed by $\mathbf{b}$'s. The frequency-domain version reconstructs the image spectrum one atom at a time, beginning with "highest frequency" atoms (whose $\mathbf{b}$'s contain the most 1's), and working towards the lower frequency atoms (whose $\mathbf{b}$'s contain fewer ones).

In particular, as will be shown, the spectrum $X^{\mathbf{b}}(\mathbf{u})$ in the atom indexed by $\mathbf{b}$ is computed via

$$X^{\mathbf{b}}(\mathbf{u}) \;=\; X_{\mathbf{b}}(\mathbf{u}) \;-\; \sum_{\mathbf{b}':\,\|\mathbf{b}'\|>\|\mathbf{b}\|} X_{\mathbf{b}}^{\mathbf{b}'}(\mathbf{u})\,, \tag{1}$$

where $X_{\mathbf{b}}(\mathbf{u})$ is the spectrum of the image samples in the bi-step lattice parametrized by $\mathbf{b}$ (a subset of the Manhattan set), the sum is over all $\mathbf{b}'$ with more ones than $\mathbf{b}$, and $X_{\mathbf{b}}^{\mathbf{b}'}(\mathbf{u})$ is the spectrum of the samples (taken with the same bi-step lattice) of the image component $x^{\mathbf{b}'}(\mathbf{t})$ corresponding to atom $\mathbf{b}'$, which has previously been reconstructed.

A discrete-space version of this requires only DFTs of the subsampling of the Manhattan samples and the previously reconstructed image components specified in the above, as well as summing and subtracting. Then an inverse DFT computes the newly reconstructed component. Summing all such components yields the reconstructed image.

The method characterized by (1), and the discrete-space version thereof, can also be carried out in the spatial domain by applying the right-hand side of (1) to the corresponding sampled images, rather than their spectra, and then applying an ideal bandpass filter that extracts just the frequency component corresponding to atom $\mathbf{b}$. The impulse responses of these filters will be given later. As will be seen, these impulse responses depend on the k_i's and λ_i's, but not the choice of bi-step lattices that comprise the Manhattan set. Moreover, the λ_i's have only a simple spatial scaling effect on the filters.

Finally, we note that the development for three dimensions benefits greatly from the efficient parametrization of bi-step lattices mentioned earlier, and that with such, it is possible to derive the M-sampling theorem and reconstruction procedure in arbitrary dimensions with essentially no additional effort or notation.

We conclude the introduction by relating the present work to previous work. Multidimensional sampling theorems, showing that images with certain spectral support regions can be reconstructed from certain samplings sets, appeared first for lattice sampling sets in Peterson and Middleton [11], then later for unions of shifted lattices, i.e., lattice cosets, [16]–[27], although they were not always described as such.

The earliest work [11], [16] required the spectral support region and sampling set to be chosen so that the spectral replicas induced by sampling did not overlap, and consequently, reconstruction could proceed simply by lowpass filtering the sampled image. For example, the approach of [16] could be used to reconstruct images from M-samples. However, it would require the images to be bandlimited to the Nyquist region of the *coarse (rectangular) lattice*, which is the intersection (rather than union) of the bi-step lattices comprising the Manhattan set.

Nonoverlapping spectral replicas were not required in later work [17]–[27], and more complex reconstruction procedures were proposed. Though not specifically intended for images, a seminal contribution stimulating a number of advances in image sampling was the multichannel, generalized sampling introduced by Papoulis [28]. For example, Papoulis' framework is broad enough to include all image sampling schemes based on lattices and unions of shifted lattices.

One difference between the present work and much past work is that we focus on a particular sampling set, namely a Manhattan set, and seek a largest possible frequency region such that any image bandlimited to such can be reconstructed from the samples. In contrast, much of the past work [17]–[19], [23], [24], [29] focused on a particular frequency support region and sought a smallest possible sampling set, constructed from lattices and shifts thereof, such that images bandlimited to this region could be reconstructed from such sampling sets. Nevertheless, some of the latter approaches could be used to reverse engineer reconstruction procedures and/or spectral support regions for Manhattan sets, as we now discuss.

One substantial line of past work applies to sampling sets that consist of a sublattice of some specified *base lattice*, together with some of its cosets, each of which is a shift of the sublattice by some base lattice point. In this case, the subsampling corresponding to each coset (including the sublattice itself) can be viewed as a *channel* in a Papoulis multichannel, generalized sampling scheme. Consequently, the method of [28] can be applied. This is the approach taken by Marks and Cheung [17]–[19]. Since a Manhattan set can be viewed as the union of what we earlier called the coarse (rectangular) lattice and some number of its cosets with respect to the *dense (rectangular) lattice*, which contains all points $\mathbf{t}$ such that for each i, its ith coordinate is an integer multiple of λ_i, the Papoulis-Marks-Cheung (PMC) approach can be applied to Manhattan sets.

In particular, Marks and Cheung focused on images with a given spectral support region and an initial base sampling lattice such that the induced spectral replicas of this support region do not overlap. They then showed that cosets of some sublattice could be removed from the base lattice until the sampling density was minimal (in the Landau sense) or approached minimal. Their method involved (a) partitioning the Nyquist region of the initial base lattice into atoms the size and shape of the Nyquist region of the sublattice, (b) counting the number of atoms of this partition that are not overlapped by any spectral replica of the designated support region induced by the initial base sampling lattice, and (c) showing that this number of sublattice cosets can be removed from the initial base lattice due to their samples being linearly dependent on other samples. If the atoms of the partition are too coarse to closely match the set of frequencies not contained in any spectral support replica, then choosing a sparser sublattice

will enable a finer partitioning, resulting in a higher fraction of the base samples being removed, which allows the sampling rate to be reduced until it equals or approaches the Landau minimum.

With hindsight, one can apply their approach to a Manhattan sampling set. For simplicity, consider a 2D case and assume $k_1 = k_2 = k$. Suppose images are bandlimited to the cross-shaped Manhattan region, and let the initial base sampling lattice and the sublattice be the dense and coarse rectangular lattices mentioned earlier. In this case, there are k^2 cosets of the sublattice (including itself). One can then see that in the partition of the Nyquist region of the base/dense lattice into atoms having the size and shape of the Nyquist region of the coarse lattice, the number of atoms that are not contained in any sampled spectra is $k^2 - (2k - 1)$. Thus, it is possible to remove all but $2k - 1$ cosets, which is precisely the number of Manhattan samples in one $k \times k$ fundamental cell of the coarse lattice. Unfortunately, the PMC approach does not determine which cosets can be removed, so it does not directly tell us if the Manhattan samples are sufficient to recover an image. While it does provide a matrix invertibility test that one can apply in any particular case to see if the Manhattan samples are sufficient, it is not clear how to analytically establish that one can remove all but the Manhattan samples in all cases. It is also not clear how the PMC approach would have lead to the discovery that the union of the Nyquist regions of the bi-step lattices is a reconstructable spectra support region for Manhattan sampling, especially in dimensions three and above. However, once it is known that the Manhattan samples are sufficient for the spectral support region found in the present paper, then the Papoulis approach will directly lead to a reconstruction algorithm.

As both the PMC and onion-peeling approaches involve partitioning frequency space, it is interesting to note that in dimension d the PMC approach requires a partition into $\prod_{i=1}^{d} k_i$ atoms, whereas the onion-peeling algorithm partitions into only 2^d atoms. The smaller size of the latter partition is due to its being closely tailored to the specific structure of Manhattan samples.

Similarly, in another line of work, Faridani [20] derived a sampling theorem and reconstruction formula for unions of shifts of one lattice. Given a spectral support region, the reconstruction involves partitioning this region in a certain way and setting up and solving a sizable number of systems of linear equations, assuming that the equations have a solution. Since a Manhattan set can be viewed as the union of shifts of a lattice (the coarse lattice) and since we know from the results of the present paper that it is possible to reconstruct M-sampled images bandlimited to the Manhattan region, one could presumably solve the resulting equations to obtain a reconstruction formula. While this is interesting, finding the partition and setting up the equations can be difficult, especially in high dimensions. Thus, as before, the onion-peeling approach proposed in this paper is more natural, intuitive and straightforward to implement.

While the PMC and Faridani approaches could be used to derive a reconstruction method for any Manhattan set, in their basic form, they do not provide direct closed form reconstruction methods, as given for example in this paper. That is, given sets of k_i's and bi-step lattices, they outline a procedure that could be followed in order to derive a reconstruction method. Then, when the k_i's or bi-step lattices are changed, the procedure must be followed again, essentially from scratch[1]. In contrast, the reconstruction methods given in the paper are closed form, requiring just step-by-step following of the reconstruction formulas, which depend explicitly on the λ_i's, k_i's and bi-step lattices. While it is conceivable that with enough work this alternative approach could be made closed form, it would appear to take much additional work, especially to make it apply to arbitrary dimensions.

Behmard [26] derived a sampling theorem and reconstruction formula for unions of shifts of more than one lattice, which includes M-sampling, as it is a more general setting than [20]. However, the *compatibility conditions* required to apply this approach are not satisfied by M-sampling and the Manhattan spectral support region.

Other work on sampling with unions of shifted lattices includes that of (a) Venkataramani and Bresler [23], [24], which considered unions of shifted lattices in one dimension, and (b) Unnikrishnan and Vetterli [27], which considered unions of shifted lattices in higher dimensions. The latter include M-sampling and a reconstruction procedure was proposed with similarities to our onion-peeling approach, but which requires the spectral support region to be convex, which rules out the Manhattan region. Indeed, one of their examples is a 2D Manhattan grid, from which images can be recovered provided their spectra are bandlimited to a circular subset of the Manhattan region. Consequently, a significantly smaller set of images is reconstructable with their procedure.

Finally, we mention that Manhattan-bandlimited spectra have been found to arise naturally in dynamic medical imaging applications, including both time-varying tomography [29] and dynamic MRI [30]. For example, Rilling et. al. [30, Fig. 1] give carotid blood velocity mapping as an example of a dynamic MRI application where a cross-shaped spectrum appears. Moreover, such spectra arise when temporal variation is localized to a small spatial area relative to the rest of the body, such as beating heart. With this motivation, Willis and Bresler [29] derived a single sampling lattice such that the cross-shaped spectral replicas did not overlap and the sampling rate was close to the Landau lower bound. In contrast, our sampling theorem also shows perfect reconstruction is possible. However, we sample with more than one lattice, the spectral replicas overlap, and the Landau bound is met exactly.

[1]The method can be derived assuming unit λ_i's and then spatially scaled for the actual λ_i's.

In summary, given that the present paper shows that images bandlimited to the union of the Nyquist regions of the bi-step lattices of a Manhattan sampling set can be perfectly reconstructed from the Manhattan samples, there are probably a number of alternative ways to derive reconstruction algorithms. In the view of the authors, the onion-peeling method, whose development was guided by the specific structure of Manhattan samples, is a natural and efficient reconstruction method with a straightforward interpretation in frequency space. It is also closed form in terms of the parameters of the Manhattan set. In addition, the union-of-bi-step-lattice viewpoint taken in this paper leads naturally to the hypothesis that the union of Nyquist regions is a support region of images that are reconstructable from Manhattan samples. It is not known if other approaches would have lead investigators to this region.

The paper is written so that the reader who is primarily interested in 2D images can focus on Sections II, III, and V.

II. Preliminaries

This section provides background and notation for sampling and lattices that will be used throughout the the paper.

Let $\mathbb{R}$ denote the real numbers, let $\mathbb{R}^d$ denote d-dimensional Euclidean space, let $\mathbb{Z}$ denote the set of all integers, and let $\mathbb{Z}^d$ denote the set of all integer-valued d-dimensional vectors. In dimension d, an image is a mapping $x(\mathbf{t}) : \mathbb{R}^d \to \mathbb{R}$, where the spatial variable is $\mathbf{t} = (t_1, \ldots, t_d)$. We restrict attention to images $x(\mathbf{t})$ that contain no delta functions or other generalized functions, and have well defined Fourier transforms containing no delta functions or other generalized functions, where by Fourier transform we mean

$$X(\mathbf{u}) \;=\; \mathcal{F}\big\{x(\mathbf{t})\big\} \;\triangleq\; \int x(\mathbf{t})\, e^{-j2\pi \mathbf{t}\cdot\mathbf{u}}\, d\mathbf{t}\,.$$

We will often refer to $X(\mathbf{u})$ as the *spectrum* of $x(\mathbf{t})$.

Sampling a d-dimensional image $x(\mathbf{t})$ means collecting its values on some countable *sampling set* $\mathbb{S}$. That is, it produces the set of values $\{x(\mathbf{t}) : \mathbf{t} \in \mathbb{S}\}$. As commonly done, one can model such sampling as multiplication of $x(\mathbf{t})$ by the *comb function* of the set $\mathbb{S}$, which produces the *sampled image*

$$x_{\mathbb{S}}(\mathbf{t}) \;\triangleq\; x(\mathbf{t})\, K_{\mathbb{S}} \sum_{\mathbf{t}' \in \mathbb{S}} \delta(\mathbf{t} - \mathbf{t}')\,, \tag{2}$$

where $K_{\mathbb{S}}$ is a normalizing constant and $\delta(\mathbf{t})$ denotes the Dirac delta function in d-space. The Fourier transform of $x_{\mathbb{S}}(\mathbf{t})$ is then called the *sampled spectrum*.

Rectangular sampling refers to sampling with a rectangular lattice. Given d and $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_d)$ with positive components, the d-dimensional *rectangular lattice* with *step vector* $\boldsymbol{\alpha}$ is a countably infinite

set of points that are spaced by integer multiples of the *step size* α_i in the ith dimension. Specifically,

$$\begin{aligned} L(\boldsymbol{\alpha}) &\triangleq \left\{\mathbf{t} : t_i \text{ is a multiple of } \alpha_i, i = 1, \ldots, d\right\} \\ &= \left\{\mathbf{t} = \mathbf{n} \odot \boldsymbol{\alpha} : \mathbf{n} \in \mathbb{Z}^d\right\}, \end{aligned}$$

where $\odot$ denotes element-wise product (also known as the Hadamard or Schur product). Alternatively, $L(\boldsymbol{\alpha})$ is the additive group generated by the basis $\alpha_1\mathbf{e}_1, \ldots, \alpha_d\mathbf{e}_d$, where $\mathbf{e}_1, \ldots, \mathbf{e}_d$ is the standard basis, i.e., $\mathbf{e}_i$ has a 1 in the ith place and 0's elsewhere. That is,

$$L(\boldsymbol{\alpha}) \triangleq \left\{\mathbf{t} = \sum_{i=1}^{d} n_i \alpha_i \mathbf{e}_i : \mathbf{n} \in \mathbb{Z}^d\right\}.$$

The *reciprocal lattice* corresponding to $L(\boldsymbol{\alpha})$ is

$$L^*(\boldsymbol{\alpha}) \triangleq L(\alpha_1^{-1}, \ldots, \alpha_d^{-1}).$$

When sampling with set $\mathbb{S} = L(\boldsymbol{\alpha})$, it is convenient to set the normalizing constant in (2) to be

$$K_{\mathbb{S}} = \textstyle\prod_{i=1}^{d} \alpha_i. \tag{3}$$

With this, the sampled image, denoted $x_{\boldsymbol{\alpha}}(\mathbf{t})$, has spectrum

$$X_{\boldsymbol{\alpha}}(\mathbf{u}) = \sum_{\mathbf{v} \in L^*(\boldsymbol{\alpha})} X(\mathbf{u} - \mathbf{v}). \tag{4}$$

¿From this, one sees that the sampled spectrum $X_{\boldsymbol{\alpha}}(\mathbf{u})$ consists of replicas of the original image spectrum $X(\mathbf{u})$, translated to the sites in frequency domain of the reciprocal lattice. The usual d-dimensional sampling theorem follows from the fact that if the support of $X(\mathbf{u})$ lies entirely within the Nyquist region[2]

$$\mathcal{N}_{\boldsymbol{\alpha}} \triangleq \left\{\mathbf{u} : |u_i| < \frac{1}{2\alpha_i},\ i = 1, \ldots, d\right\},$$

then said replicas do not overlap, and consequently, the original spectrum can be recovered by extracting the portion of the sampled image spectrum in the Nyquist region.

III. Two-Dimensional Manhattan Sampling

As introduced earlier and depicted in Fig. 1(a), Manhattan sampling (M-sampling) uses locations spaced closely along a grid of horizontal and vertical lines. In particular, we assume there is a sample at the origin, as well as samples spaced λ_1 apart on horizontal lines spaced $k_2\lambda_2$ apart, and samples spaced λ_2

[2]In this paper, script variables such as $\mathcal{N}, \mathcal{B}$ or $\mathcal{A}$ will usually denote subsets of frequency space.

apart on vertical lines spaced $k_1\lambda_1$ apart, where $\lambda_i > 0$ and k_1, k_2 are integers greater than one[3]. The issue, now, is to find an as large as possible set of images that can be perfectly reconstructed from these samples, as well as an efficient procedure for doing so.

A first thought is to model M-sampling as multiplying the given image $x(\mathbf{t})$ by a comb function having delta functions at the Manhattan sampling locations, and then to analyze the spectra of the resulting sampled image. Since this comb function has the same periodicity as a comb function for the *coarse lattice* $L_C \triangleq L(k_1\lambda_1, k_2\lambda_2)$, the replicas of the image spectrum lie at frequency sites in the reciprocal lattice L_C^*, or a subset thereof. Thus, perfect reconstruction is possible for images bandlimited to the Nyquist region $\mathcal{N}_C$ of the coarse lattice L_C. However, since such reconstructions need only use samples in the coarse lattice, it may be that a larger set of images is reconstructable from the full Manhattan grid. On the other hand, if images are bandlimited to a region larger than $\mathcal{N}_C$, e.g., a scaling of the Nyquist region such as $(1+\epsilon)\mathcal{N}_C$, $\epsilon > 0$, then the spectral replicas induced by an M-sampling comb may overlap. Even if this does not eliminate the possibility of perfect reconstruction, it will at least complicate the analysis.

Accordingly, we pursue an approach that does not rely on nonoverlapping replicas, but derives from the key observation that the Manhattan sampling set can be viewed as the *union* of two rectangular lattices. Let us initially focus on what can be recovered from the samples of each lattice by itself. The *horizontal lattice*, $L_H \triangleq L(\lambda_1, k_2\lambda_2)$, densely samples in the horizontal direction and coarsely samples in the vertical direction; the *vertical lattice*, $L_V \triangleq L(k_1\lambda_1, \lambda_2)$, coarsely samples in the horizontal direction and densely samples in the vertical direction; and the sampling set for M-sampling is

$$M(\boldsymbol{\lambda}; \mathbf{k}) \;=\; L_H \cup L_V \,.$$

Note also that the intersection of the two lattices is the coarse lattice L_C, whose comb function has the same periodicity as a comb function for the Manhattan grid.

Clearly, all images bandlimited to the Nyquist region $\mathcal{N}_H$ of the horizontal lattice L_H can be recovered from just the samples in this lattice. Likewise, all images bandlimited to the Nyquist region $\mathcal{N}_V$ of the vertical lattice L_V can be recovered from just the samples in this lattice. Each of these by itself leads to a larger recoverable set of images than the set recoverable from sampling with the coarse lattice L_C. However, neither type of sampling and reconstruction uses all of the M-samples.

We now show how images bandlimited to the union of the Nyquist regions of the horizontal and vertical lattices can be recovered from the full set of M-samples. Specifically, suppose image $x(\mathbf{t})$ is

[3]We require $k_1, k_2 > 1$ since if $k_1 = 1$ or $k_2 = 1$, the sampling set reduces to a normal rectangular lattice.

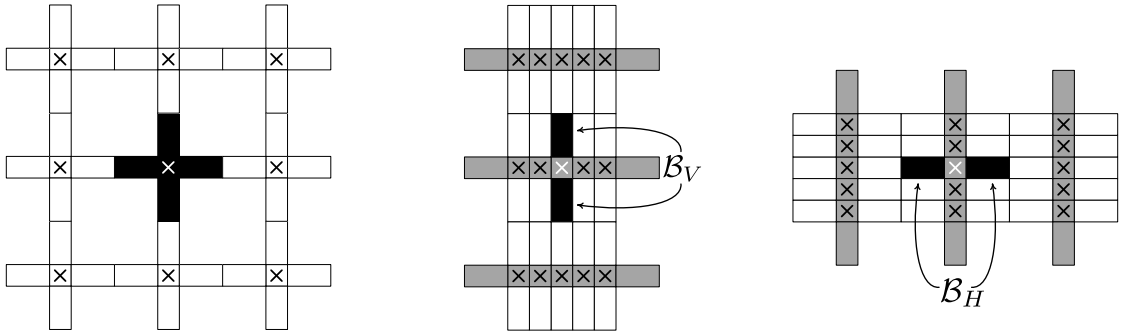

Fig. 2. For Manhattan sampling with $\lambda_1 = \lambda_2$ and $k_1 = k_2 = 3$: (a) Support of the sampled spectrum for an image bandlimited to the Manhattan region, when sampled with the fine lattice L_{λ_1,λ_2}. The original spectrum is black with a white $\times$ in its center, whereas replicas are white with a black $\times$ in their centers. (b) Support of the sampled spectrum when sampled with the vertical lattice $L_{k_1\lambda_1,\lambda_2}$. Gray indicates regions where replicas overlap either the original spectrum or each other. (c) Same as (b), except that the sampling is with the horizontal lattice $L_{\lambda_1,k_2\lambda_2}$.

bandlimited to $\mathcal{M}(\boldsymbol{\lambda};\mathbf{k}) = \mathcal{N}_H \cup \mathcal{N}_V$, which is the cross-shaped region shown in Fig. 1(c). First, consider only the samples of $x(\mathbf{t})$ taken on the vertical lattice. Since the cross-shaped region $\mathcal{M}(\boldsymbol{\lambda};\mathbf{k})$ is not contained in the Nyquist region $\mathcal{N}_V$, the replicas of $X(\mathbf{u})$ may overlap in the spectrum of the vertically sampled image, as illustrated in Fig. 2(b), However, certain portions of each cross-shaped replica cannot be overlapped, and thus these portions of the spectra of $x(\mathbf{t})$ can be immediately recovered.

Specifically, it is easy to see that with vertical sampling, the vertical highpass region $\mathcal{B}_V \triangleq \mathcal{N}_V - \mathcal{N}_C$ is not overlapped. Thus, with $I_D(\mathbf{u})$ denoting the indicator function of some set D and $X^V(\mathbf{u}) \triangleq X(\mathbf{u})I_{\mathcal{B}_V}(\mathbf{u})$ denoting the portion of $X(\mathbf{u})$ in $\mathcal{B}_V$, one sees that from the vertical samples and their spectrum $X_V(\mathbf{u})$, one can immediately recover $X^V(\mathbf{u})$ via $X^V(\mathbf{u}) = X_V(\mathbf{u})I_{\mathcal{B}_V}(\mathbf{u})$. Likewise from the horizontal samples and their spectrum $X_H(\mathbf{u})$, the horizontal highpass region $\mathcal{B}_H \triangleq \mathcal{N}_H - \mathcal{N}_C$ is not overlapped. Thus, one can immediately recover $X^H(\mathbf{u}) \triangleq X(\mathbf{u})I_{\mathcal{B}_H}(\mathbf{u}) = X_H(\mathbf{u})I_{\mathcal{B}_H}(\mathbf{u})$.[4]

Since $X^V(\mathbf{u})$ and $X^H(\mathbf{u})$ are now known, and $X(\mathbf{u})$ is bandlimited to $\mathcal{M}(\boldsymbol{\lambda},\mathbf{k}) = \mathcal{B}_H \cup \mathcal{B}_V \cup \mathcal{N}_C$, it remains only to find $X^C(\mathbf{u}) \triangleq X(\mathbf{u})I_{\mathcal{N}_C}(\mathbf{u})$. It will then follow that $X(\mathbf{u}) = X^V(\mathbf{u}) + X^H(\mathbf{u}) + X^C(\mathbf{u})$. Inverse transforms will give $x(\mathbf{t}) = x^V(\mathbf{t}) + x^H(\mathbf{t}) + x^C(\mathbf{t})$.

To determine $X^C(\mathbf{u})$, consider the vertical sampling of $x(\mathbf{t})$, and observe in Fig. 2(b) that the overlap of the image spectrum $X(\mathbf{u})$ in $\mathcal{N}_C$ by the various spectral replicas is due only to replications of the

[4]Throughout the paper, a superscript on an image x or spectrum X will usually pertain to a frequency region, and a subscript will usually pertain to a sampling.

horizontal highpass frequency components in $\mathcal{B}_H$. Since these have already been determined, it ought to be possible subtract their effects.

To see that this can be done, let us focus on $X_V(\mathbf{u})\, I_{\mathcal{N}_C}(\mathbf{u})$. ¿From (4) and the fact that $X(\mathbf{u}) = 0$ for $\mathbf{u} \notin \mathcal{M}(\boldsymbol{\lambda}, \mathbf{k})$, we have

$$X_V(\mathbf{u})\, I_{\mathcal{N}_C}(\mathbf{u}) \;=\; \sum_{i=-n}^{n} X\Big(u_1 - \frac{i}{k_1\lambda_1}, u_2\Big)\, I_{\mathcal{N}_C}(\mathbf{u})\,,$$

where $n = \big\lfloor \frac{k_1}{2} \big\rfloor$. Now using $X(\mathbf{u}) = X^V(\mathbf{u}) + X^H(\mathbf{u}) + X^C(\mathbf{u})$ in the above along with the facts that

(a) $X^V\big(u_1 - \frac{i}{k_1\lambda_1}, u_2\big)\; I_{\mathcal{N}_C}(\mathbf{u}) = 0$ for all i,

(b) $X^H\big(u_1 - \frac{i}{k_1\lambda_1}, u_2\big)\, I_{\mathcal{N}_C}(\mathbf{u}) = 0$ for $i = 0$,

(c) $X^C\big(u_1 - \frac{i}{k_1\lambda_1}, u_2\big)\, I_{\mathcal{N}_C}(\mathbf{u}) = 0$ unless $i = 0$,

we find

$$X_V(\mathbf{u})\, I_{\mathcal{N}_C}(\mathbf{u}) \;=\; X^C(\mathbf{u}) \,+\, Y(\mathbf{u})\, I_{\mathcal{N}_C}(\mathbf{u})\,, \tag{5}$$

where

$$\begin{aligned} Y(\mathbf{u}) \;&\triangleq\; \sum_{\substack{i=-n \\ i\neq 0}}^{n} X^H\Big(u_1 - \frac{i}{k_1\lambda_1}, u_2\Big) \\ &= \sum_{\substack{i=-n \\ i\neq 0}}^{n} X_H\Big(u_1 - \frac{i}{k_1\lambda_1}, u_2\Big) I_{\mathcal{B}_H}\Big(u_1 - \frac{i}{k_1\lambda_1}, u_2\Big). \end{aligned} \tag{6}$$

where the last equality uses the fact, mentioned earlier, that $X^H(\mathbf{u}) = X_H(\mathbf{u}) I_{\mathcal{B}_H}(\mathbf{u})$. Notice that $Y(\mathbf{u})$ is the component of $X_V(\mathbf{u})$ due to aliasing by replicas of $X^H(\mathbf{u})$, and is directly computable from the horizontal samples. It follows from (5) that $X^C(\mathbf{u}) = (X_V(\mathbf{u}) - Y(\mathbf{u}))\, I_{\mathcal{N}_C}(\mathbf{u})$.

In summary, a procedure for recovering a cross bandlimited x from its M-samples is

1) Compute the spectra, $X_H(\mathbf{u})$ and $X_V(\mathbf{u})$, of the horizontally and vertically dense samples, respectively.
2) From $X_H(\mathbf{u})$, compute $Y(\mathbf{u})$ for $\mathbf{u} \in \mathcal{N}_C$.
3) Let

$$\widehat{X}(\mathbf{u}) = \begin{cases} X_V(\mathbf{u}), & \mathbf{u} \in \mathcal{B}_V \\ X_H(\mathbf{u}), & \mathbf{u} \in \mathcal{B}_H \\ X_V(\mathbf{u}) - Y(\mathbf{u}), & \mathbf{u} \in \mathcal{N}_C \end{cases}. \tag{7}$$

4) Let $\widehat{x}(\mathbf{t})$ be the inverse Fourier transform of $\widehat{X}(\mathbf{u})$.

This result is summarized in the following.

Theorem 1: 2D Manhattan sampling theorem. Given $\lambda_1, \lambda_2 > 0$ and integers k_1, k_2 greater than 1, any image $x(\mathbf{t})$ whose Fourier transform is bandlimited to the cross-shaped region $\mathcal{M}(\boldsymbol{\lambda}; \mathbf{k})$ can be recovered from its M-samples in $M(\boldsymbol{\lambda}; \mathbf{k})$ with the procedure given above.

The following alternative expression for $X_V(\mathbf{u})\, I_{N_C}(\mathbf{u})$ will lead to an easier to implement procedure for discrete-space images with finite support (presented later). Using (4) and $X(\mathbf{u}) = X^V(\mathbf{u}) + X^H(\mathbf{u}) + X^C(\mathbf{u})$, we find

$$\begin{aligned} X_V(\mathbf{u})\, I_{\mathcal{N}_C}(\mathbf{u}) &= \sum_{\mathbf{v}\in L_V^*} \Big(X^V(\mathbf{u}-\mathbf{v}) + X^H(\mathbf{u}-\mathbf{v}) + X^C(\mathbf{u}-\mathbf{v}) \Big) I_{\mathcal{N}_C}(\mathbf{u}) \\ &= X^C(\mathbf{u}) \,+\, \sum_{\mathbf{v}\in L_V^*} X^H(\mathbf{u}-\mathbf{v}) \; I_{\mathcal{N}_C}(\mathbf{u}) \\ &= X^C(\mathbf{u}) \,+\, Y'(\mathbf{u}) I_{\mathcal{N}_C}(\mathbf{u})\,, \end{aligned}$$

where

$$Y'(\mathbf{u}) \;\triangleq\; \sum_{\mathbf{v}\in L_V^*} X^H(\mathbf{u}-\mathbf{v})\,. \tag{8}$$

It follows that $Y(\mathbf{u})$ in the procedure given previously can be replaced by $Y'(\mathbf{u})$. The advantage is that, as shown below, $Y'(\mathbf{u})$ can be computed with Fourier transforms instead of a summation. To show this, let $\mathcal{S}_V$ denote the vertical sampling operator, which when applied to an image $z(\mathbf{t})$ produces $z_{L_V}(t)$ as defined by (2). We recognize the summation in (8) as the sampled spectrum when the image $x^H(\mathbf{t})$ is vertically sampled. Since $x^H(\mathbf{t})$, and consequently $X^H(\mathbf{u})$, can be computed from the horizontal samples,

$$\begin{aligned} Y'(\mathbf{u}) \;&=\; \mathcal{F}\big\{\mathcal{S}_V\big\{x^H(\mathbf{t})\big\}\big\} \;=\; \mathcal{F}\big\{\mathcal{S}_V\big\{\mathcal{F}^{-1}\big\{X^H(\mathbf{u})\big\}\big\}\big\} \\ &=\; \mathcal{F}\Big\{\mathcal{S}_V\Big\{\mathcal{F}^{-1}\big\{I_{\mathcal{B}_H}(\mathbf{u})\mathcal{F}\big\{x_H(\mathbf{t})\big\}\big\}\Big\}\Big\}\,. \end{aligned}$$

While the above may initially appear complex[5], in the discrete-space, finite-support case discussed shortly, it will lead to a simple procedure that avoids the summations in (6) and (8).

Maximality, in the Landau sense, of the set of reconstructable images

The sampling density of an $M(\boldsymbol{\lambda}; \mathbf{k})$ M-sampling set is

$$\rho = \frac{k_1 + k_2 - 1}{k_1 k_2 \lambda_1 \lambda_2}\,,$$

[5]Note also that the expression (8) for $Y'(\mathbf{u})$ contains more terms in the sum than the corresponding expression (6) for $Y(\mathbf{u})$.

since any $k_1\lambda_1 \times k_2\lambda_2$ rectangle in $\mathbb{R}^2$ contains $k_1 + k_2 - 1$ samples. In the frequency domain, the area of the Manhattan-bandlimited region, denoted $|\mathcal{M}(\boldsymbol{\lambda}; \mathbf{k})|$, is the sum of the areas of $\mathcal{B}_H$, $\mathcal{B}_V$ and $\mathcal{N}_C$. Alternatively, it is sum of the areas of the Nyquist regions corresponding to the horizontal and vertical sampling lattices, minus the area of their intersection. Either way, this may be written as

$$|\mathcal{B}_H| + |\mathcal{B}_V| + |\mathcal{N}_C| = \frac{1}{k_1\lambda_1\lambda_2} + \frac{1}{k_2\lambda_1\lambda_2} - \frac{1}{k_1k_2\lambda_1\lambda_2}\,,$$

which simplifies to the previous expression for sampling density ρ. Thus, the set of images bandlimited to the Manhattan region $\mathcal{M}(\boldsymbol{\lambda}; \mathbf{k})$ is a maximal set of reconstructable images in the Landau sense for the M-sampling grid $M(\boldsymbol{\lambda}; \mathbf{k})$.

Discrete-space images

In this section, we briefly consider M-sampling of discrete-space images. Such images might be created by rectangularly sampling a continuous-space image, or they might exist only as discrete-space objects. In any case, we consider an image to be a mapping $x[\mathbf{t}] : \mathbb{T} \to \mathbb{R}$ where $\mathbb{T}$ is either the (infinite) integer lattice $\mathbb{Z}^2$, or a finite subset of the form $\mathbb{T} = \{\mathbf{t} : 0 \le t_1 \le T_1 - 1, 0 \le t_2 \le T_2 - 1\}$ for some positive integers T_1, T_2. *Sampling* $x[\mathbf{t}]$ refers to collecting a subset of its values. In the infinite support case, $\mathbb{T} = \mathbb{Z}^2$, a Manhattan grid $M(\boldsymbol{\lambda}, \mathbf{k})$ is once again defined to be the union of a horizontal lattice $L_H \triangleq L(\lambda_1, k_2\lambda_2)$ and a vertical lattice $L_V \triangleq L(k_1\lambda_1, \lambda_2)$, except that now each lattice must be a sublattice of the integer lattice $\mathbb{Z}^d$, i.e., λ_1, λ_2 must be positive integers. In this case, we assume the discrete-space Fourier transform of $x[\mathbf{t}]$ is well defined and contains no delta functions or other generalized functions. In the finite support case, a Manhattan grid is formed in a similar way, namely, $M(\boldsymbol{\lambda}, \mathbf{k}) = L_H \cup L_V$, where now L_H and L_V are truncated to the finite $\mathbb{T}$.

(a) Infinite-support discrete-space images: In this case, Theorem 1 holds with only trivial changes, as does the reconstruction procedure. Specifically, the only required changes are: (i) replace the continuous-space Fourier transform as the formula for a spectrum with the discrete-space Fourier transform, and (ii) scale all specified frequencies by 2π, such as those defining Nyquist regions and $\mathcal{M}(\boldsymbol{\lambda}; \mathbf{k})$.

(b) Finite-support discrete-space images: In this case, as is customary, we use the Discrete Fourier Transform (DFT) as the formula for the spectrum of an image:

$$X[\mathbf{u}] \;= \sum_{\mathbf{t}\in\mathbb{T}} x[\mathbf{t}]\, e^{-j2\pi(\frac{u_1}{T_1}t_1 + \frac{u_2}{T_2}t_2)}, \quad \mathbf{u} \in \mathbb{T}\,.$$

The conventional sampling theorem (c.f. [26]) for discrete-space images with spatial support $\mathbb{T}$ (defined by T_1, T_2) sampled with a rectangular lattice $L(\alpha_1, \alpha_2)$ limited to $\mathbb{T}$ says that an image $x[\mathbf{t}]$ with support

$\mathbb{T}$ can be recovered from its samples in this lattice if T_1 and T_2 are integer multiples of α_1 and α_2, respectively, and its DFT $X[\mathbf{u}]$ is zero outside the (discrete) Nyquist region

$$\widetilde{\mathcal{N}}_{\alpha_1,\alpha_2} \triangleq \Big\{\mathbf{u} \in \mathbb{T} : \text{ for } i = 1\,\&\,2,\ 0 \le u_i < \frac{T_i}{2\alpha_i} \text{ or } T_i - \frac{T_i}{2\alpha_i} < u_i \le T_i - 1\Big\}.$$

Now suppose a finite-support discrete-space image $x[\mathbf{t}]$ is sampled on the Manhattan grid $M(\boldsymbol{\lambda}, \mathbf{k})$ and is bandlimited to the cross-shaped *Manhattan region*

$$\widetilde{\mathcal{M}}(\boldsymbol{\lambda}, \mathbf{k}) \triangleq \widetilde{\mathcal{N}}_H \cup \widetilde{\mathcal{N}}_V\,,$$

where $\widetilde{\mathcal{N}}_H$ and $\widetilde{\mathcal{N}}_V$ are the Nyquist regions of the horizontal and vertical lattices, respectively. Assuming T_1 and T_2 are integer multiples of $k_1\lambda_1$ and $k_2\lambda_2$, respectively, a straightforward adaptation of the analysis for continuous-space images shows that from the samples in the vertical lattice L_V, one can recover the spectrum $X[\mathbf{u}]$ in the highpass region $\widetilde{\mathcal{B}}_V \triangleq \widetilde{\mathcal{N}}_V - \widetilde{\mathcal{N}}_C$, where $\widetilde{\mathcal{N}}_V$ and $\widetilde{\mathcal{N}}_C$ are the Nyquist regions of the vertical and coarse lattices, respectively. Specifically, from the spectrum $X_V[\mathbf{u}]$ of the vertically sampled image $x_V[\mathbf{t}]$ (with scaling as in (2)-(3)), one recovers $X^V[\mathbf{u}] \triangleq X[\mathbf{u}]I_{\widetilde{\mathcal{B}}_V}[\mathbf{u}] = X_V[\mathbf{u}]I_{\widetilde{\mathcal{B}}_V}[\mathbf{u}]$. Likewise, from the samples in the horizontal lattice L_H, one can recover the spectrum in the highpass region $\widetilde{\mathcal{B}}_H \triangleq \widetilde{\mathcal{N}}_H - \widetilde{\mathcal{N}}_C$ from the spectrum $X_H[\mathbf{u}]$ of the horizontally sampled image $x_H[\mathbf{t}]$ via $X^H[\mathbf{u}] \triangleq X[\mathbf{u}]I_{\widetilde{\mathcal{B}}_H}(\mathbf{u}) = X_H[\mathbf{u}]I_{\widetilde{\mathcal{B}}_H}[\mathbf{u}]$. Finally, the spectrum in the Nyquist region $\widetilde{\mathcal{N}}_C$ of the coarse lattice can be determined via $X^C[\mathbf{u}] = (X_V[\mathbf{u}] - Y[\mathbf{u}])\, I_{\widetilde{\mathcal{N}}_C}[\mathbf{u}]$, where

$$Y[\mathbf{u}] \;=\; k_1\lambda_1\lambda_2 \sum_{r=1}^{k_1\lambda_1 - 1} X^H\Big[\Big(u_1 - r\frac{T_1}{k_1\lambda_1}\Big) \bmod T_1, u_2\Big].$$

This leads to the following.

Theorem 2: 2D discrete-space, finite-support Manhattan sampling theorem. If T_1 and T_2 are integer multiples of $k_1\lambda_1$ and $k_2\lambda_2$, respectively, then an image $x[\mathbf{t}]$ with finite support $\mathbb{T}$ can be recovered from its M-samples in $M(\boldsymbol{\lambda}, \mathbf{k})$ if its DFT $X[\mathbf{u}]$ is zero outside the Manhattan region $\widetilde{\mathcal{M}}(\boldsymbol{\lambda}, \mathbf{k})$.

Reconstruction procedure:

Given the samples in Manhattan grid $M(\boldsymbol{\lambda}, \mathbf{k})$ of an image $x[\mathbf{t}]$ bandlimited to $\widetilde{\mathcal{M}}(\boldsymbol{\lambda}, \mathbf{k})$, the following adaptation of the continuous-space procedure recovers the entire $x[\mathbf{t}]$.

1) Let $x_V[\mathbf{t}]$ equal $k_1\lambda_1\lambda_2\, x[\mathbf{t}]$ on the vertical lattice L_V and zero otherwise, and let $x_H[\mathbf{t}]$ equal $k_2\lambda_1\lambda_2\, x[\mathbf{t}]$ on the horizontal lattice L_H and zero otherwise. Compute $X_V[\mathbf{u}] = \text{DFT}\big\{x_V[\mathbf{t}]\big\}$ and $X_H[\mathbf{u}] = \text{DFT}\big\{x_H[\mathbf{t}]\big\}$.
2) Compute the "alias subtraction" term

$$Y'[\mathbf{u}] \;=\; \text{DFT}\Big\{\widetilde{\mathcal{S}}_V\Big\{\text{IDFT}\big\{I_{\widetilde{\mathcal{B}}_H}[\mathbf{u}]\,\text{DFT}\{x_H[\mathbf{t}]\}\big\}\Big\}\Big\},$$

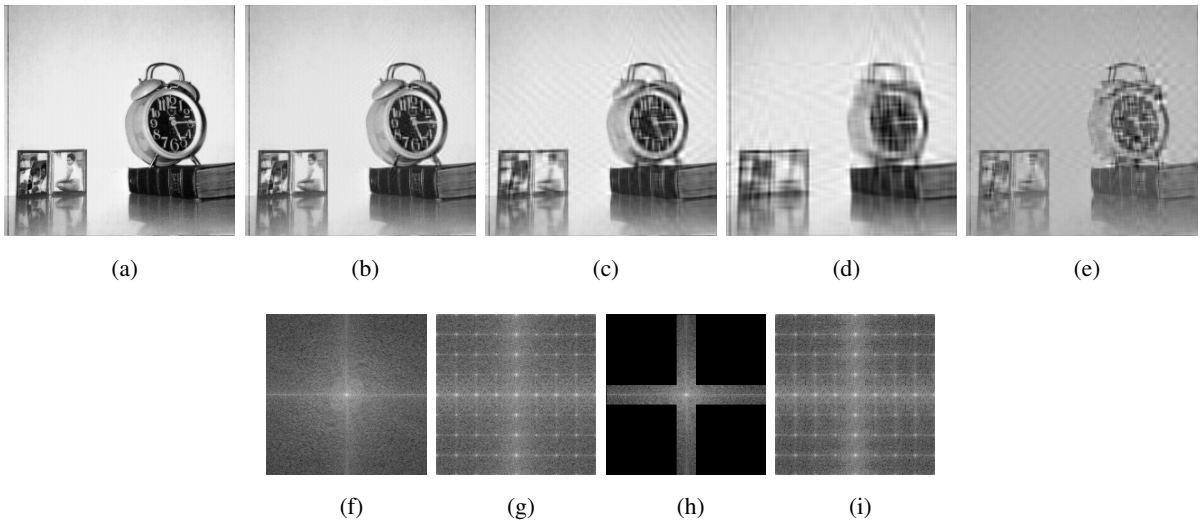

(a) (b) (c) (d) (e)

(f) (g) (h) (i)

Fig. 3. (a) Original 256×256 image. (b) Image bandlimited to Manhattan region $\widetilde{\mathcal{M}}(\boldsymbol{\lambda}, \mathbf{k})$, with $k_1 = k_2 = 4$ and $\lambda_1 = \lambda_2 = 1$. (c) Same as (b) except $k_1 = k_2 = 8$. (d) Same as (c) except $\lambda_1 = \lambda_2 = 2$. (Note: after spectra were zeroed outside $\widetilde{\mathcal{M}}(\boldsymbol{\lambda}, \mathbf{k})$, inverse transforms were applied, negligible imaginary parts were discarded, and images were quantized to $\{0, 1, \ldots, 255\}$.) (e) Image sampled with parameters of (c) and reconstructed without first bandlimiting to Manhattan region. Log magnitude spectra: (f) original image; (g) original image sampled with parameters of (c); (h) original image bandlimited with parameters of (c); (i) bandlimited image (c) sampled with parameters of (c);

where $\widetilde{\mathcal{S}}_V$ denotes the vertical sampling operator that, when applied to an image $z[\mathbf{t}]$, produces an image that is $k_1\lambda_1\lambda_2\, z[\mathbf{t}]$ on L_V, and zero elsewhere.

3) Compute the spectrum:

$$\widehat{X}[\mathbf{u}] \;=\; \begin{cases} X_V[\mathbf{u}], & \mathbf{u} \in \widetilde{\mathcal{B}}_V \\ X_H[\mathbf{u}], & \mathbf{u} \in \widetilde{\mathcal{B}}_H \\ X_V[\mathbf{u}] - Y'[\mathbf{u}], & \mathbf{u} \in \widetilde{N}_C \end{cases} .$$

4) Invert the spectrum:

$$x[\mathbf{t}] \;=\; \mathrm{IDFT}\big\{X[\mathbf{u}]\big\}\,.$$

This reconstruction procedure uses 5 DFT/IDFT operations, each requiring $\mathcal{O}(N \log N)$ arithmetic operations when implemented with an FFT, where $N = T_1 T_2$, plus three pairwise additions of $T_1 \times T_2$ matrices, each requiring $T_1 T_2$ additions, plus instances of setting matrix elements to zero. In summary, the complexity of reconstruction, which is dominated by the FFT's, is $\mathcal{O}(N \log N)$ operations per image.

Note that few real world, finite-support images will satisfy the conditions of Theorem 2. As a result, to apply M-sampling to a real world image, the image can be pre-processed by zero-padding so that

its dimensions are multiples of $k_1\lambda_1$ and $k_2\lambda_2$, respectively, and "Manhattan filtering" by taking the DFT and setting to zero all coefficients outside of $\widetilde{\mathcal{M}}(\boldsymbol{\lambda}, \mathbf{k})$. Such padded and filtered images can be recovered perfectly from their M-samples. To illustrate the effects of such filtering, which heavily suppresses diagonal frequencies, Fig. 3 shows a finite-support image and its filtering with several choices of parameters. It also shows the spectra of the image before and after bandlimiting, the spectrum of the sampled image, with and without bandlimiting, and the effect of sampling and reconstruction without first pre-filtering. Note that the image was chosen to have sharp edges surrounded by a smooth background in order that one can easily see the ringing due to bandlimiting.

IV. Higher-Dimensional Manhattan Sampling

A. Introduction

As mentioned earlier, in any dimension $d \geq 3$ there are a number of possible d-dimensional Manhattan sets. Each is a finite union of rectangular lattices, each defined by step sizes that in dimension i are constrained to be λ_i or $k_i\lambda_i$, where each λ_i is a positive constant called the *dense spacing* in dimension i, and each k_i is an integer greater than 1 called the *sampling factor* in dimension i. Such a rectangular lattice will be called a $(\boldsymbol{\lambda}, \mathbf{k})$-lattice, where $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_d)$ is its *dense spacing vector* and $\mathbf{k} = (k_1, \ldots, k_d)$ is its *sampling factor vector*. It will also be called a $(d, \boldsymbol{\lambda}, \mathbf{k})$-lattice when we wish to emphasize d, and a *bi-step lattice* when we do not wish to specify parameters. (The term "bi-step" emphasizes that each step size α_i can only take one of two values: λ_i or $k_i\lambda_i$). Accordingly, to specify a d-dimensional *Manhattan set*, one specifies a dense spacing vector $\boldsymbol{\lambda}$, a sampling factor vector $\mathbf{k}$, and a collection of $(\boldsymbol{\lambda}, \mathbf{k})$-lattices.

To efficiently characterize a $(\boldsymbol{\lambda}, \mathbf{k})$-lattice, we let $\mathbf{b} = (b_1, \ldots, b_d)$ denote a d-dimensional vector, called its *bi-step indicator vector*, or more concisely *bi-step vector*, that indicates the dimensions along which the bi-step lattice is dense, according to the convention $b_i = 1$ if the step size is λ_i in dimension i and 0 if the step size is $k_i\lambda_i$. Thus, the $(\boldsymbol{\lambda}, \mathbf{k})$-lattice specified by $\mathbf{b}$ is $L_{\boldsymbol{\lambda},\mathbf{k},\mathbf{b}} \triangleq L(\boldsymbol{\alpha}_{\mathbf{b}})$, with $\boldsymbol{\alpha}_{\mathbf{b}} = (\alpha_{\mathbf{b},1}, \ldots, \alpha_{\mathbf{b},d})$ defined by

$$\alpha_{\mathbf{b},i} \;=\; \begin{cases} \lambda_i, & b_i = 1 \\ k_i\lambda_i, & b_i = 0 \end{cases} . \tag{9}$$

or equivalently,

$$L_{\boldsymbol{\lambda},\mathbf{k},\mathbf{b}} \;\triangleq\; \left\{\mathbf{t} : t_i \text{ is a multiple of } k_i\lambda_i \text{ for } i \text{ s.t. } b_i = 0, \text{ and a multiple of } \lambda_i \text{ for other } i\right\}. \tag{10}$$

Note that we generally consider d, $\boldsymbol{\lambda}$ and $\mathbf{k}$ to be fixed, and so as an abbreviation and slight abuse of notation, we usually write $L_{\mathbf{b}}$ instead of $L_{\boldsymbol{\lambda},\mathbf{k},\mathbf{b}}$. It will also be useful to let $x_{\mathbf{b}}(\mathbf{t})$ and $X_{\mathbf{b}}(\mathbf{u})$ denote, respectively, the sampled image and the sampled spectrum due to sampling $x(\mathbf{t})$ with $L_{\mathbf{b}}$.

The following summarizes.

Definition 1: Given dimension d, dense spacing vector $\boldsymbol{\lambda}$, sampling factor vector $\mathbf{k}$ (all of its components are integers greater than 1), and a finite collection of $(d, \boldsymbol{\lambda}, \mathbf{k})$-lattices specified by the bi-step vectors in $B = \{\mathbf{b}_1, \ldots \mathbf{b}_m\}$, the corresponding $(d, \boldsymbol{\lambda}, \mathbf{k}, B)$-*Manhattan (sampling) set* is

$$M\big(d, \boldsymbol{\lambda}, \mathbf{k}, B\big) \triangleq \bigcup_{j=1}^{m} L_{\mathbf{b}_j} . \tag{11}$$

As $d, \boldsymbol{\lambda}$ and $\mathbf{k}$ will be considered fixed, we usually write $M\big(B\big)$ instead of $M\big(d, \boldsymbol{\lambda}, \mathbf{k}, B\big)$. B will be called a *Manhattan collection* or *M-collection* for short.

B. Examples and properties of bi-step lattices

It is useful to call attention to certain bi-step lattices. One is the *dense lattice* $L_{\mathbf{1}}$ corresponding to the bi-step vector $\mathbf{b} = \mathbf{1} \triangleq (1, \ldots, 1)$. It is a rectangular lattice with step vector $\boldsymbol{\lambda}$ that contains every other $(\boldsymbol{\lambda}, \mathbf{k})$-lattice. Another is the *coarse lattice* $L_{\mathbf{0}}$ corresponding to $\mathbf{b} = \mathbf{0} \triangleq (0, \ldots, 0)$, which is the rectangular lattice with step vector $\boldsymbol{\alpha} = \mathbf{k} \odot \boldsymbol{\lambda}$ and which is contained in every other $(\boldsymbol{\lambda}, \mathbf{k})$-lattice. As mentioned in the introduction for 3D Manhattan sets, it will be useful to consider the partition of $\mathbb{R}^d$ induced by the coarse lattice, whose cells are $k_1\lambda_1 \times \ldots \times k_d\lambda_d$ orthotopes (hyper-rectangles) with corners at lattice points. These orthotopes will be called *fundamental cells*. The coarse lattice contains just the corners of these fundamental cells; other $(\boldsymbol{\lambda}, \mathbf{k})$-lattices may contain points on their edges and faces, but only the dense lattice $L_{\mathbf{1}}$ contains points in their interiors.

A third lattice to consider is $L_{\mathbf{e}_i}$ corresponding to bi-step vector $\mathbf{b} = \mathbf{e}_i$, which can be viewed as a collection of points spaced densely on lines parallel to $\mathbf{e}_i$, with one line passing through each point of $(d-1)$-dimensional cubic lattice $L(k_1\lambda_1, \ldots, k_{i-1}\lambda_{i-1}, k_{i+1}\lambda_{i+1}, \ldots, k_d\lambda_d)$. Finally, we mention the lattice corresponding to $\mathbf{b} = \mathbf{1} - \mathbf{e}_i$, which can be viewed as sampling densely on shifts of the $d-1$ dimensional lattice $L(\lambda_1, \ldots, \lambda_{i-1}, \lambda_{i+1}, \ldots, \lambda_d)$ spaced $k_i\lambda_i$ apart.

Given two (binary) bi-step vectors $\mathbf{b}_1$ and $\mathbf{b}_2$, define their *union* $\mathbf{b}_1 \vee \mathbf{b}_2$ and *intersection* $\mathbf{b}_1 \wedge \mathbf{b}_2$ to be their element-wise 'OR' and 'AND', respectively, and define $\mathbf{b}_1 \subset \mathbf{b}_2$ to mean $\mathbf{b}_1 \wedge \mathbf{b}_2 = \mathbf{b}_1$. Define the complement to be $\mathbf{b}^c \triangleq \mathbf{1} - \mathbf{b}$, and the *Hamming weight* or simply *weight* $\|\mathbf{b}\|$ to be the number of ones contained in $\mathbf{b}$.

The following are useful properties of γ representations of bi-step lattices.

Fact 1: Considering $(d, \boldsymbol{\lambda}, \mathbf{k})$-lattices,

(a) $L_{\mathbf{b}_1} \subset L_{\mathbf{b}_2}$ if and only if $\mathbf{b}_1 \subset \mathbf{b}_2$,

(b) $L_{\mathbf{b}_1} = L_{\mathbf{b}_2}$ if and only if $\mathbf{b}_1 = \mathbf{b}_2$,

(c) $L_{\mathbf{b}_1} \cap L_{\mathbf{b}_2} \;=\; L_{\mathbf{b}_1 \wedge \mathbf{b}_2}$,

(d) If $L_{\widetilde{\mathbf{b}}} \subset \bigcup_{j=1}^m L_{\mathbf{b}_j}$, then for some j, $L_{\widetilde{\mathbf{b}}} \subset L_{\mathbf{b}_j}$, and consequently from (a), $\widetilde{\mathbf{b}} \subset \mathbf{b}_j$.

Proof:

(a) and (b) are elementary.

(c) $L_{\mathbf{b}_1} \cap L_{\mathbf{b}_2}$

$= \big\{\mathbf{t} : t_i$ is a multiple of $k_i\lambda_i$ for all i s.t. $b_{1,i} = 0$ or $b_{2,i} = 0$, and t_i is a multiple of λ_i for other $i\big\}$

$= \big\{\mathbf{t} : t_i$ is a multiple of $k_i\lambda_i$ for all i s.t. $(\mathbf{b}_1 \wedge \mathbf{b}_2)_i = 0$ and t_i is a multiple of λ_i for other $i\big\}$

$= L_{\mathbf{b}_1 \wedge \mathbf{b}_2}$

(d) As is well known, for $m = 2$ this property derives from just the group nature of lattices, but not for larger values of m. Accordingly, to prove it for arbitrary m, we need to use properties of $(\boldsymbol{\lambda}, \mathbf{k})$ lattices. Specifically, we demonstrate the contrapositive. Suppose $L_{\widetilde{\mathbf{b}}} \not\subset L_{\mathbf{b}_j}$, $j = 1, \ldots, m$. Then from (a), for each $1 \le j \le m$, $\widetilde{\mathbf{b}} \not\subset \mathbf{b}_j$, and so there exists i_j such that $\widetilde{b}_{i_j} = 1$ and $b_{j,i_j} = 0$. Let I denote the set of all such i_j's, and let $\mathbf{t} = (t_1, \ldots, t_d)$ be defined by

$$t_i = \begin{cases} (k_i + 1)\lambda_i, & \text{if } i \in I \\ k_i\lambda_i, & \text{otherwise} \end{cases} .$$

Referring to (10), we see that $\mathbf{t} \in L_{\widetilde{\mathbf{b}}}$, because the only dimensions i for which t_i is not a multiple of $k_i\lambda_i$ are those in I, in which case $\widetilde{b}_i = 1$, i.e., the lattice $L_{\widetilde{\mathbf{b}}}$ is dense in dimension i. Moreover, again referring to (10), we see that for each $j \in \{1, \ldots, m\}$, $\mathbf{t} \notin L_{\mathbf{b}_j}$, because $t_{i_j} = (k_{i_j} + 1)\lambda_i$ is not a multiple of $k_{i_j}\lambda_i$, yet $b_{j,i_j} = 0$, i.e., the lattice $L_{\mathbf{b}_j}$ is coarse in dimension i_j. It follows that $\mathbf{t} \notin \cup_{j=1}^m L_{\mathbf{b}_j}$. Hence, $L_{\widetilde{\mathbf{b}}} \not\subset \cup_{j=1}^m L_{\mathbf{b}_j}$. □

Among other things, (b) shows there is a one-to-one correspondence between bi-step vectors $\mathbf{b}$ and $(d, \boldsymbol{\lambda}, \mathbf{k})$ lattices. which verifies that the $\mathbf{b}$'s are valid representations of the $(\boldsymbol{\lambda}, \mathbf{k})$-lattices. Also, since there are 2^d possible bi-step vectors $\mathbf{b}$, it follows that there are 2^d distinct $(d, \boldsymbol{\lambda}, \mathbf{k})$-lattices.

Note that while (c) shows that the intersection of two $(d, \boldsymbol{\lambda}, \mathbf{k})$-lattices is another $(d, \boldsymbol{\lambda}, \mathbf{k})$-lattice, such is not true for the union, which is why a Manhattan set is not ordinarily a lattice. For example, when

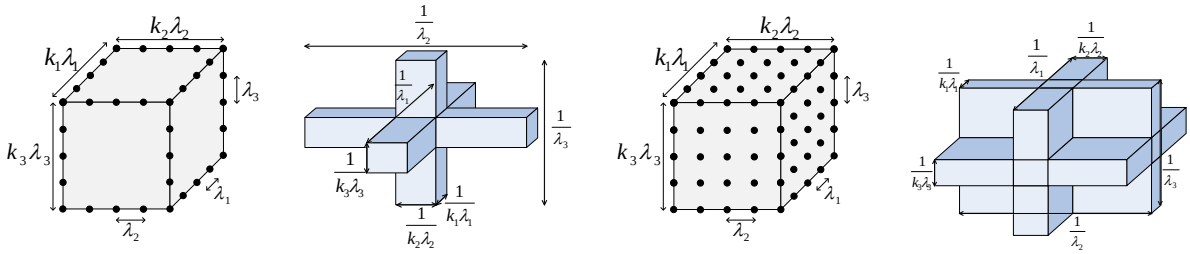

(a) (b) (c) (d)

Fig. 4. Examples of 3D Manhattan sampling $M(B)$ and their corresponding Manhattan regions $\mathcal{M}(B)$. (a) Manhattan lines $B = \{(1,0,0),(0,1,0),(0,0,1)\}$, and (b) its corresponding Manhattan region. (c) Manhattan facets $B = \{(1,1,0),(1,0,1),(0,1,1)\}$, and (d) its corresponding Manhattan region.

$d = 2$, $L_{\mathbf{e}_1} \cup L_{\mathbf{e}_2}$ is the Manhattan set $M\big(\{\mathbf{e}_1, \mathbf{e}_2\}\big)$ shown in Fig. 1(a), which is not a lattice and does not equal $L_{\mathbf{e}_1 \vee \mathbf{e}_2}$, which is the dense lattice $L_{(1,1)}$.

C. Examples of Manhattan sets

Several types of Manhattan sets deserve special attention.

1) *Manhattan lines* is a Manhattan set $M(B)$ specified by $B = \{\mathbf{e}_1, \ldots, \mathbf{e}_d\}$. In this case the samples are taken on the 1D edges of the fundamental cells, i.e., on d orthogonal sets of parallel lines in $\mathbb{R}^d$. See Figures 1(a) and 4(a) for illustrations of Manhattan lines in two and three dimensions, respectively.
2) *Manhattan facets* is a Manhattan set $M(B)$ specified by $B = \{\mathbf{e}_1^c, \ldots, \mathbf{e}_d^c\}$. Sampling on a set of Manhattan facets is analogous to sampling densely along d orthogonal sets of parallel hyperplanes in $\mathbb{R}^d$. See Figures 1(a) and 4(c) for illustrations of Manhattan facets in two and three dimensions, respectively. For $d = 2$, Manhattan facets and lines are identical.
3) Though technically any $(\boldsymbol{\lambda}, \mathbf{k})$-lattice, including the coarse and dense lattices, is a Manhattan set, we focus on Manhattan sets that are not lattices, which we call *proper*.
4) Video sampling: Let $d = 3$, let $i = 1, 2$ be spatial dimensions and let $i = 3$ be a temporal dimension. Whereas video is most commonly sampled with a rectangular lattice, say $L(\lambda_1, \lambda_2, \lambda_3)$, other samplings are possible, for example, the Manhattan set $M(3, \boldsymbol{\lambda}, \mathbf{k}, B)$ specified by $B = \{\mathbf{e}_3^c, \mathbf{e}_3\}$ uses fine spatial sampling every $k_3\lambda_3$ seconds and spatial subsampling with factors k_1 and k_2 at times that are other multiples of λ_3 seconds.

D. Alternate representations of Manhattan sets

By the definition of a Manhattan set (11) and Fact 1(a), augmenting an M-collection B by a subset $\mathbf{b}'$ of some $\mathbf{b}$ in B does not change the resulting Manhattan set. Thus many M-collections can generate the same Manhattan set. There are, however, unique largest and smallest M-collections that generate any given Manhattan set. To find these, we make use of the following.

Fact 2: Let B and B' be M-collections. Then,

(a) $M(B) \subset M(B')$ if $B \subset B'$.

(b) $M(B) \subset M(B')$ if and only if for each $\mathbf{b} \in B$ there is $\mathbf{b}' \in B'$ such that $L_{\mathbf{b}} \subset L_{\mathbf{b}'}$, or equivalently by Fact 1(a), $\mathbf{b} \subset \mathbf{b}'$.

(c) It is not true that $M(B) = M(B')$ implies $B = B'$, or that $M(B) \subset M(B')$ implies $B \subset B'$.

Proof:

(a) is obvious.

(b) If for each $\mathbf{b} \in B$ there is $\mathbf{b}' \in B'$ s.t. $L_{\mathbf{b}} \subset L_{\mathbf{b}'}$, then $M(B) = \cup_{\mathbf{b} \in B} L_{\mathbf{b}} \subset \cup_{\mathbf{b}' \in B'} L_{\mathbf{b}'} = M(B')$. Conversely, if $M(B) \subset M(B')$, then for each $\mathbf{b} \in B$, $L_{\mathbf{b}} \subset M(B') = \cup_{b' \in B'} L_{\mathbf{b}'}$, and Fact 1(d) implies $L_{\mathbf{b}} \subset L_{\mathbf{b}'}$ for some $\mathbf{b}' \in B'$.

(c) Part (a) shows that adding to B some subset of some $\mathbf{b} \in B$ that is not already in B yields $B' \neq B$ such that $M(B) = M(B')$. In this same case, $M(B') \subset M(B)$, but $B' \not\subset B$. □

The unique largest M-collection that generates $M(B)$ is

$$\overline{B} \triangleq \left\{\mathbf{b}' : \mathbf{b}' \subset \mathbf{b} \text{ for some } \mathbf{b} \in B\right\},$$

which we call the *closure* of B. Fact 2(a) implies $M(B) \subset M(\overline{B}) = M(B)$, and Fact 2(b) implies $M(\overline{B}) \subset M(B)$. Hence $M(\overline{B}) = M(B)$. $M(\overline{B})$ is the largest M-collection generating $M(B)$ because if $M(B') = M(B)$, and $\mathbf{b}' \in B'$, then Fact 2(b) implies there is a $\mathbf{b} \in B$ such that $L_{\mathbf{b}'} \subset L_{\mathbf{b}}$, and this implies $\mathbf{b}' \in \overline{B}$. Hence, $B' \subset \overline{B}$.

Removing all elements of an M-collection B that are subsets of another element results in the unique smallest M-collection generating $M(B)$, which we denote $\underline{B}$.

E. Manhattan sampling density

The *density* of a Manhattan set $M(B)$, i.e., the number of samples per unit area, is obviously less than the sum of the densities of the lattices of which it is the union, because each lattice contains all points in

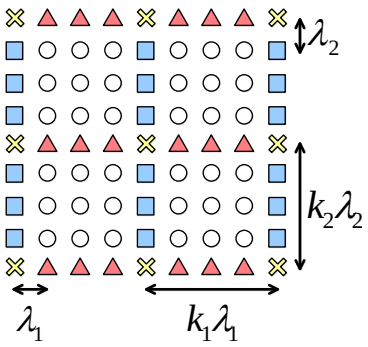

Fig. 5. Partitioning of 2D Manhattan grid $M(\{\mathbf{e}_1, \mathbf{e}_2\})$ with sampling factors $k_1 = k_2 = 4$, which is the union of the yellow $\times$'s in $V_{(0,0)}$, the red $\triangle$'s in $V_{(1,0)}$, and the blue $\square$'s in $V_{(0,2)}$. The white $\circ$'s are in $V_{(1,1)}$, which is disjoint from this Manhattan grid.

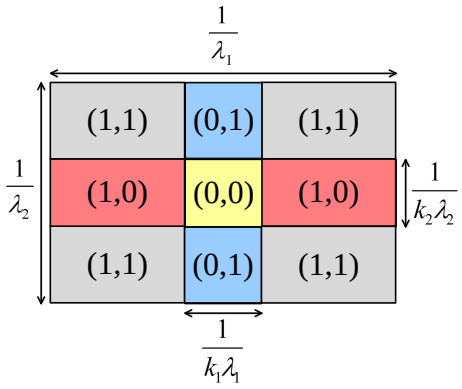

Fig. 6. M-partition of $\mathcal{N}_{\mathbb{D}}$ for $d = 2$, $k_1 = 5$, $k_2 = 3$. Frequency $\mathbf{u} = 0$ lies at the center. Each M-atom $A^{\mathbf{b}}$ is identified by its $\mathbf{b}$. Note that the cross-shaped Manhattan region $\mathcal{M}(\{\mathbf{e}_1, \mathbf{e}_2\})$ is also partitioned by M-atoms; in particular, $\mathcal{M}(\{\mathbf{e}_1, \mathbf{e}_2\}) = A^{(0,0)} \cup A^{(1,0)} \cup A^{(0,1)}$.

the coarse lattice $L_{\mathbf{0}}$. Accordingly, we partition the dense lattice $L_{\mathbf{1}}$ in such a way that for any B, $M(B)$ is the union of atoms of this partition, and its density can be computed by summing their densities.

To obtain a suitable partition, let us group into one atom all sites $\mathbf{t}$ of $L_{\mathbf{1}}$ having the same answers to the following d questions – "Is t_i not a multiple of $\lambda_i k_i$?" – for $i = 1, \ldots, d$. Specifically, with a binary vector $\mathbf{b} = (b_1, \ldots, b_d)$ indicating the set of i's for which the answers are "yes", consider the partition $\{V_{\mathbf{b}} : \mathbf{b} \subset \{0,1\}^d\}$, where the atom corresponding to $\mathbf{b}$ is

$$V_{\mathbf{b}} \triangleq \{\mathbf{t} : \ t_i \text{ is a multiple of } k_i\lambda_i \text{ for } i \text{ s.t. } b_i = 0,$$
$$\text{and } t_i \text{ is a multiple of } \lambda_i, \text{ but not } k_i\lambda_i, \text{ for } i \text{ s.t. } b_i = 1\}$$

Fig. 5 illustrates the partitioning of a 2D Manhattan grid. Essentially, it is a partition of the dense lattice into collections of cosets of the coarse lattice.

It is clear from the definition that no $\mathbf{t}$ can lie in both $V_{\mathbf{b}}$ and $V_{\mathbf{b}'}$ for $\mathbf{b} \neq \mathbf{b}'$. Hence, the $V_{\mathbf{b}}$'s are disjoint. By comparing the above to the definition (10) of $L_{\mathbf{b}}$, one sees that $V_{\mathbf{b}'} \subset L_{\mathbf{b}}$ if and only if $\mathbf{b}' \subset \mathbf{b}$. It follows that for any $\mathbf{b}$, $\cup_{\mathbf{b}' \subset b} V_{\mathbf{b}'} \subset L_{\mathbf{b}}$. Conversely, if $\mathbf{t} \in L_{\mathbf{b}}$, then it is easily seen that $\mathbf{t} \in L_{\mathbf{b}'}$ for $\mathbf{b}'$ defined by $b'_i = 0$ for i such that t_i is a multiple of $\lambda_i k_i$ and $b'_i = 1$ otherwise. It follows that $\cup_{\mathbf{b}' \subset b} V_{\mathbf{b}'} = L_{\mathbf{b}}$, i.e., $\{V_{\mathbf{b}}\}$ partitions any bi-step lattice, including $L_{\mathbf{1}}$. Moreover, since $M(B) = \cup_{\mathbf{b} \in B} L_{\mathbf{b}}$, one can also write $M(B) = \cup_{\mathbf{b} \in \overline{B}} V_{\mathbf{b}}$, i.e., $\{V_{\mathbf{b}}\}$ partitions any Manhattan grid.

With this partition in mind, the density of the Manhattan set $M(B)$ is now obtained by summing the

densities of each $V_{\mathbf{b}}$, $\mathbf{b} \in \overline{B}$. Consider the points of $V_{\mathbf{b}}$ in the fundamental cell

$$F_{\mathbf{k},\boldsymbol{\lambda}} \triangleq \bigtimes_{i=1}^{d} [0, k_i\lambda_i)\,,$$

which has a corner at the origin, lies entirely in the positive hyper-quadrant, and has volume is $\prod_{i=1}^{d} k_i\lambda_i$. We see that $V_{\mathbf{b}} \cap F_{\mathbf{k},\boldsymbol{\lambda}}$ is the Cartesian product of d sets $A_1, \ldots, A_d$, where $A_i = \{\lambda_i, 2\lambda_i, \ldots, (k_i-1)\lambda_i\}$ if $b_i = 1$ and $A_i = \{0\}$ if $b_i = 0$. Since $V_{\mathbf{b}} \cap F_{\mathbf{k},\boldsymbol{\lambda}}$ contains $\prod_{i:b_i=1}(k_i - 1)$ points, and since the density of $V_{\mathbf{b}}$ is this number divided by the volume of $F_{\mathbf{k},\boldsymbol{\lambda}}$, the density of $M(B)$ is

$$\rho(B) \;=\; \frac{\sum_{\mathbf{b}\in\overline{B}} \prod_{i:b_i=1}(k_i - 1)}{\prod_{i=1}^{d} k_i} \times \frac{1}{\prod_{i=1}^{d} \lambda_i}\,. \tag{12}$$

For example, the densities of several Manhattan sets in three dimensions are given in Table I.

Manhattan set	Γ	$\rho(\Gamma)$
Manhattan lines	$\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$	$(3k-2)/k^3$
Video sampling example	$\{\mathbf{e}_3^c, \mathbf{e}_3\}$	$(k^2+k-1)/k^3$
Manhattan facets	$\{\mathbf{e}_1^c, \mathbf{e}_2^c, \mathbf{e}_3^c\}$	$(3k^2-3k+1)/k^3$

TABLE I

DENSITY OF SEVERAL 3-DIMENSIONAL MANHATTAN SETS WITH $k_i = k$ AND $\lambda_i = 1$ FOR ALL i.

F. Manhattan partition of frequency space

As mentioned earlier, our approach to reconstructing an appropriately bandlimited image $x(\mathbf{t})$ from M-samples $M(B)$ involves sequentially reconstructing regions of its spectrum $X(\mathbf{u})$. Specifically, each region will be recovered from the samples in some collection of bi-step lattices contained in the Manhattan set. This section describes a partition of frequency space, some of whose atoms will be the reconstructable regions.

Let $\mathcal{N}_{\mathbf{b}}$ denote the Nyquist region for bi-step lattice $L_{\mathbf{b}}$, i.e.,

$$\mathcal{N}_{\mathbf{b}} \triangleq \left\{\mathbf{u} : |u_i| < \frac{1}{2\alpha_{\mathbf{b},i}},\, i = 1, \ldots, d\right\},$$

with step sizes $\alpha_{\mathbf{b},i}$ given by (9). For future reference we note that

$$\mathcal{N}_{\mathbf{b}'} \subset \mathcal{N}_{\mathbf{b}} \text{ if and only if } \mathbf{b}' \subset \mathbf{b}\,. \tag{13}$$

Since any $(d, \boldsymbol{\lambda}, \mathbf{k})$ Manhattan set is a subset of the dense lattice $L_{\mathbf{1}}$, it follows that the appropriate bandlimitation for images reconstructable from any such Manhattan set or any $(d, \boldsymbol{\lambda}, \mathbf{k})$ lattice is a subset of the Nyquist region of the dense lattice, namely,

$$\mathcal{N}_{\mathbf{1}} \;=\; \left\{\mathbf{u} : |\mathbf{u}_i| < \frac{1}{2\lambda_i},\, i = 1, \ldots, d\right\}.$$

Thus, we need only partition $\mathcal{N}_{\mathbf{1}}$.

Definition 2: The *Manhattan partition* (*M-partition*) of $\mathcal{N}_{\mathbf{1}}$ is $\left\{\mathcal{A}^{\mathbf{b}} : \mathbf{b} \in \{0,1\}^d\right\}$ where $\mathcal{A}^{\mathbf{b}}$ is the *Manhattan atom*[6] *(M-atom)*

$$\mathcal{A}^{\mathbf{b}} \;\triangleq\; a_1^{\mathbf{b}} \times \cdots \times a_d^{\mathbf{b}}\,,$$

and $a_i^{\mathbf{b}}$ is the interval, or union of two intervals,

$$a_i^{\mathbf{b}} \triangleq \begin{cases} \left\{u_i : \frac{1}{2k_i\lambda_i} \le |u_i| < \frac{1}{2\lambda_i}\right)\}, & b_i = 1 \\ \left\{u_i : |u_i| < \frac{1}{2k_i\lambda_i}\right\}, & b_i = 0 \end{cases}.$$

Thus, $\mathcal{A}^{\mathbf{b}}$ is highpass for all dimensions such that $b_i = 1$ and lowpass for all other dimensions. The *weight* of atom $\mathcal{A}^{\mathbf{b}}$ is $\|\mathbf{b}\|$, the weight of $\mathbf{b}$.

The M-partition is illustrated in Fig. 6 in the case of $d = 2$, $k_1 = 5$, $k_2 = 3$. We now make several easy to deduce, but important, observations.

1) M-atom $\mathcal{A}^{\mathbf{b}}$ is the Cartesian product of lowpass intervals $(-\frac{1}{2k_i\lambda_i}, \frac{1}{2k_i\lambda_i})$ along each dimension such that $b_i = 0$, and of the union of two highpass intervals $(-\frac{1}{2\lambda_i}, -\frac{1}{2k_i\lambda_i}] \cup [\frac{1}{2k_i\lambda_i}, \frac{1}{2\lambda_i})$ along each dimension such that $b_i = 1$. Thus, $\mathcal{A}^{\mathbf{b}}$ is the union of $2^{\|\mathbf{b}\|}$ disjoint orthotopes in $\mathbb{R}^d$.
2) The M-atoms are disjoint, and their union is $\mathcal{N}_{\mathbf{1}}$. Hence, they comprise a partition of $\mathcal{N}_{\mathbf{1}}$.
3) The M-atom $\mathcal{A}^{\mathbf{b}}$ is a subset of the Nyquist region $\mathcal{N}_{\mathbf{b}}$. Equality holds only for $\mathbf{b} = \mathbf{0}$.
4) The weight $\|\mathbf{b}\|$ of an atom $\mathcal{A}^{\mathbf{b}}$ is a rough indicator of how highpass or lowpass is the atom.
5) The lowest weight M-atom, $\mathcal{A}^{\mathbf{0}}$, is lowpass in all dimensions and equals the Nyquist region $\mathcal{N}_{\mathbf{0}}$ of the coarse lattice $L_{\mathbf{0}}$, which is the bi-step lattice with smallest and lowest frequency Nyquist region.
6) The highest weight M-atom, $\mathcal{A}^{\mathbf{1}}$, is highpass in all dimensions and contains the highpass "corners" of $\mathcal{N}_{\mathbf{1}}$. Its volume is at least large as that of any other atom, and usually larger. We will see later

[6]Consistent with previous conventions, $\mathbf{b}$ is a superscript because it determines a frequency region (an M-atom $\mathcal{A}^{\mathbf{b}}$), whereas it is a subscript when specifying the Nyquist region $\mathcal{N}_{\mathbf{b}}$ of a bi-step lattice $L_{\mathbf{b}}$, precisely because it prescribes a sampling. Except for $\mathbf{b} = \mathbf{0}$, the M-atoms are highpass regions, whereas Nyquist regions are lowpass.

that no proper $(d, \boldsymbol{\lambda}, \mathbf{k})$ Manhattan set permits the reconstruction of these corners since they can only be recovered by sampling densely along every dimension, i.e., they are only recoverable if we sample on the dense lattice $L_{\mathbf{1}}$.

7) If an image $x(\mathbf{t})$ is bandlimited to $\mathcal{N}_{\mathbf{1}}$, then both $x(\mathbf{t})$ and its spectrum $X(\mathbf{u})$ can be decomposed into sums of M-atom components:

$$x(\mathbf{t}) = \sum_{\mathbf{b}\in\{0,1\}^d} x^{\mathbf{b}}(\mathbf{t})$$

$$X(\mathbf{u}) = \sum_{\mathbf{b}\in\{0,1\}^d} X^{\mathbf{b}}(\mathbf{u})\,, \tag{14}$$

where $X^{\mathbf{b}}(\mathbf{u}) = X(\mathbf{u})$ for $\mathbf{u} \in \mathcal{A}^{\mathbf{b}}$, $X^{\mathbf{b}}(\mathbf{u}) = 0$ otherwise, and $x^{\mathbf{b}}(\mathbf{t})$ is the inverse transform of $X^{\mathbf{b}}(\mathbf{u})$. We shall refer to $x^{b}(\mathbf{t})$ and $X^{\mathbf{b}}(\mathbf{u})$ as *Manhattan atoms*, or simply *atoms*, of $x(\mathbf{t})$ and $X(\mathbf{u})$, respectively.

It will also be important that the M-atoms can partition the Nyquist region $\mathcal{N}_{\mathbf{b}}$ corresponding to any bi-step lattice $L_{\mathbf{b}}$, as shown below.

Fact 3:

(a) For any $\mathbf{b} \in \{0,1\}^d$, the $2^{\|\mathbf{b}\|}$ M-atoms in $\{\mathcal{A}^{\mathbf{b}'} : \mathbf{b}' \subset \mathbf{b}\}$ partition $\mathcal{N}_{\mathbf{b}}$ in the sense that $\mathcal{N}_{\mathbf{b}} = \cup_{\mathbf{b}'\subset\mathbf{b}}\mathcal{A}^{\mathbf{b}'}$.

(b) $\mathbf{b}' \subset \mathbf{b}$ if and only if $\mathcal{A}^{\mathbf{b}'} \subset \mathcal{N}_{\mathbf{b}}$.

Proof:

(a) Since the elements of $\{\mathcal{A}^{\mathbf{b}'} : \mathbf{b}' \subset \mathbf{b}\}$ are disjoint and contained in $\mathcal{N}_{\mathbf{b}}$, it suffices to show

$$\mathcal{N}_{\mathbf{b}} \;\subset\; \bigcup_{\mathbf{b}'\subset\mathbf{b}} \mathcal{A}^{\mathbf{b}'}\,. \tag{15}$$

Accordingly, suppose $\mathbf{u} \in \mathcal{N}_{\mathbf{b}}$. It is then easy to see that $\mathbf{u} \in \mathcal{A}^{\mathbf{b}'}$, where

$$b_i' = \begin{cases} 1, & \frac{1}{2k_i\lambda_i} \le |u_i| < \frac{1}{2\lambda_i} \\ 0, & |u_i| \le \frac{1}{2k_i\lambda_i} \end{cases},$$

which demonstrates (15).

(b) First, if $\mathbf{b}' \subset \mathbf{b}$, then by (13) and the third observation after the definition of M-atom, $\mathcal{A}^{\mathbf{b}'} \subset \mathcal{N}_{\mathbf{b}'} \subset \mathcal{N}_{\mathbf{b}}$. Conversely, if $\mathcal{A}^{\mathbf{b}'} \subset \mathcal{N}_{\mathbf{b}}$, then by part (a) $\mathcal{A}^{\mathbf{b}'}$ must be one of the atoms whose union is $\mathcal{N}_{\mathbf{b}}$. Hence, $\mathbf{b}' \subset \mathbf{b}$. □

G. Spectral replication induced by bi-step lattice sampling

For a Manhattan set $M(B)$, the reconstruction algorithm to follow will reconstruct an image $x(\mathbf{t})$ by reconstructing its spectrum $X(\mathbf{u})$ one Manhattan atom at a time, in an order to be specified later. In particular, for each $\mathbf{b} \in B$, it will reconstruct atom $X^{\mathbf{b}}(\mathbf{u})$ from the subset of samples corresponding to the bi-step lattice $L_{\mathbf{b}}$, taking into account the aliasing due to previously reconstructed atoms. Using the more suggestive $\mathbf{s} = (s_1, \ldots, s_d)$, rather than $\mathbf{b}$, to denote a bi-step vector that characterizes a sampling, then as reviewed in Section II, sampling with $L_{\mathbf{s}}$ replicates the spectrum $X(\mathbf{u})$ at all sites in the reciprocal lattice $L_{\mathbf{s}}^*$. Moreover, if $X(\mathbf{u})$ is bandlimited to $\mathcal{N}_{\mathbf{1}}$, substituting (14) into (4) gives the following decomposition of the sampled spectrum:

$$X_{\mathbf{s}}(\mathbf{u}) \;=\; \sum_{\mathbf{v} \in L_{\mathbf{s}}^*} \sum_{\mathbf{b} \in \{0,1\}^d} X^{\mathbf{b}}(\mathbf{u} - \mathbf{v})\,. \tag{16}$$

We will refer to the term $X^{\mathbf{b}}(\mathbf{u}-\mathbf{v})$, and its spectral support $\mathcal{A}^{\mathbf{b}}+\mathbf{v} \triangleq \{\mathbf{u}+\mathbf{v} : \mathbf{u} \in \mathcal{A}^{\mathbf{b}}\}$, as the replica of atom $X^{\mathbf{b}}(\mathbf{u})$, respectively, $\mathcal{A}^{\mathbf{b}}$, at site $\mathbf{v}$. Using this terminology, one sees that reconstructing atom $X^{\mathbf{s}}(\mathbf{u})$ from the sampled spectrum $X_{\mathbf{s}}(\mathbf{u})$ requires accounting for the potential aliasing, i.e., overlap, of the replicas of the various atoms of $X(\mathbf{u})$ on $X^{\mathbf{s}}(\mathbf{u})$. This requires knowing which replicas of each atom will alias, i.e. overlap, $\mathcal{A}^{\mathbf{s}}$. More specifically, since the algorithm will only apply to images whose spectral support is limited to some subset of the Manhattan atoms, for any pair of bi-step vectors $\mathbf{b}$ and $\mathbf{b}'$, we will need to know whether sampling with bi-step lattice $L_{\mathbf{s}}$ causes a replica of atom $\mathcal{A}^{\mathbf{b}'}$ (at some site $\mathbf{v} \in L_{\mathbf{s}}^*$) to overlap atom $\mathcal{A}^{\mathbf{b}}$ of the original spectrum.

Such overlap questions are answered by the following lemma and its corollary. Let $R_{\mathbf{s}}^{\mathbf{b}'}$ denote the union of the replicas of all $\mathbf{b}'$ atoms induced by sampling with $L_{\mathbf{s}}$. That is,

$$R_{\mathbf{s}}^{\mathbf{b}'} \;\triangleq\; \bigcup_{\mathbf{v} \in L_{\mathbf{s}}^* - \mathbf{0}} \left[\mathcal{A}^{\mathbf{b}'} + \mathbf{v}\right].$$

Lemma 1: Consider sampling with $L_{\mathbf{s}}$.

(a) For all $\mathbf{b}, \mathbf{b}' \subset \mathbf{s}$, no replica of $\mathcal{A}^{\mathbf{b}'}$ overlaps $\mathcal{A}^{\mathbf{b}}$, i.e., $R_{\mathbf{s}}^{\mathbf{b}'} \cap \mathcal{A}^{\mathbf{b}} = \emptyset$.

(b) The replicas of $\mathcal{A}^{\mathbf{b}'}$ induced by sampling with $L_{\mathbf{s}}$ do not overlap $\mathcal{A}^{\mathbf{b}}$ if there exists at least one dimension i such that $s_i = 1$ and $b_i \neq b_i'$. That is, $R_{\mathbf{s}}^{\mathbf{b}'} \cap \mathcal{A}^{\mathbf{b}} = \emptyset$ if $(\mathbf{b} \oplus \mathbf{b}') \wedge \mathbf{s} \neq \mathbf{0}$, where $\mathbf{b} \oplus \mathbf{b}'$ denotes element-wise exclusive or (XOR).

Proof:

(a) If $\mathbf{b}, \mathbf{b}' \subset \mathbf{s}$, then Fact 3(b) shows $\mathcal{A}^{\mathbf{b}}, \mathcal{A}^{\mathbf{b}'} \subset \mathcal{N}_{\mathbf{s}}$. Since the sampling theorem for conventional rectangular lattice sampling shows that replicas of $\mathcal{N}_{\mathbf{s}}$ do not overlap $\mathcal{N}_{\mathbf{s}}$, it follows that the replicas of $\mathcal{A}^{\mathbf{b}'}$ cannot overlap $\mathcal{A}^{\mathbf{b}}$.

(b) Let us compare the M-atom

$$\mathcal{A}^{\mathbf{b}} \;=\; a_1^{\mathbf{b}} \times \cdots \times a_d^{\mathbf{b}}$$

to an arbitrary replica in $R_{\mathbf{s}}^{\mathbf{b}'}$:

$$\mathcal{A}^{\mathbf{b}'} + \mathbf{v} \;=\; (a_1^{\mathbf{b}'} + v_1) \times \cdots \times (a_d^{\mathbf{b}'} + v_d)\,,$$

for $\mathbf{v} \in L_{\mathbf{s}}^* - \{\mathbf{0}\}$. Note that $\mathcal{A}^{\mathbf{b}}$ and $\mathcal{A}^{\mathbf{b}'} + \mathbf{v}$ are disjoint if and only if $a_i^{\mathbf{b}} \cap (a_i^{\mathbf{b}'} + v_i) = \emptyset$, for some i.

If, as hypothesized in the lemma, $(\mathbf{b} \oplus \mathbf{b}') \wedge \mathbf{s} \neq \mathbf{0}$. Then there must exist i such that $s_i = 1$ and either $b_i = 1$, $b_i' = 0$ or $b_i = 0$, $b_i' = 1$. First, consider the case that $b_i = 1$, $b_i' = 0$. Then

$$a_i^{\mathbf{b}} \;=\; \Big\{u_i : \frac{1}{2k_i\lambda_i} \leq |u_i| < \frac{1}{2\lambda_i}\Big\}$$

and

$$a_i^{\mathbf{b}'} + v_i \;=\; \Big\{u_i : |u_i| < \frac{1}{2k_i\lambda_i}\Big\} + v_i\,.$$

Since $s_i = 1$, we have $v_i = \frac{n_i}{\lambda_i}$ for some n_i, and one sees from the above that no matter the value of n_i, $(a_i^{\mathbf{b}'} + v_i) \cap a_i^{\mathbf{b}} = \emptyset$. Hence, $(\mathcal{A}^{\mathbf{b}'} + \mathbf{v}) \cap \mathcal{A}^{\mathbf{b}} = \emptyset$, and so $R_{\mathbf{s}}^{\mathbf{b}'} \cap \mathcal{A}^{\mathbf{b}} = \emptyset$. A similar argument applies for the case that $b_i = 0$, $b_i' = 1$. □

The following will provide a key step in showing how to reconstruct appropriately bandlimited images.

Corollary 1: If $\|\mathbf{b}'\| \leq \|\mathbf{s}\|$, then replicas in $R_{\mathbf{s}}^{\mathbf{b}'}$ do not overlap $\mathcal{A}^{\mathbf{s}}$.

Proof: We will apply Lemma 1 with $\mathbf{b} = \mathbf{s}$. If $\mathbf{b}' = \mathbf{s}$, then Part (a) of Lemma 1 shows $R_{\mathbf{s}}^{\mathbf{b}'} \cap \mathcal{A}^{\mathbf{s}} = \emptyset$. If $\mathbf{b}' \neq \mathbf{s}$ and $\|\mathbf{b}'\| \leq \|\mathbf{s}\|$, then there must exist i such that $s_i = 1$ and $b_i' = 0$. Therefore, $(\mathbf{s} \oplus \mathbf{b}') \wedge \mathbf{s} \neq \mathbf{0}$, and Part (b) of Lemma 1 shows $R_{\mathbf{s}}^{\mathbf{b}'} \cap \mathcal{A}^{\mathbf{s}} = \emptyset$. □

H. The multidimensional Manhattan sampling theorem

Given a Manhattan set $M(B)$, consider its *Manhattan region*, which is defined to be the union of the Nyquist regions of the bi-step lattices of which it is the union:

$$\mathcal{M}(B) \;\triangleq\; \bigcup_{\mathbf{b} \in B} \mathcal{N}_{\mathbf{b}} \;=\; \bigcup_{\mathbf{b} \in B} \bigcup_{\mathbf{b}' \subset \mathbf{b}} \mathcal{A}^{\mathbf{b}'} \;=\; \bigcup_{\mathbf{b}' \in \overline{B}} A^{\mathbf{b}'},$$

where the second equality uses Fact 3(a). In this section we show that images bandlimited to $\mathcal{M}(B)$ can be recovered from their M-samples in $M(B)$; we give an explicit procedure for reconstructing such

images from their samples in $M(B)$; and we show that the set of such bandlimited images is maximal in the Landau sense.

The key steps are the next two lemmas. The first shows that for any image $x(\mathbf{t})$ whose spectrum $X(\mathbf{u})$ is bandlimited to $\mathcal{M}(B)$, the portion of $X(\mathbf{u})$ in any highest weight M-atom $\mathcal{A}^{\mathbf{b}}$ can be easily recovered from the samples in $L_{\mathbf{b}}$, which are a subset of the M-samples. Specifically, $X^{\mathbf{b}}(\mathbf{u})$ can be recovered simply by extracting the $\mathcal{A}^{\mathbf{b}}$ portion of the sampled spectrum $X_{\mathbf{b}}(\mathbf{u})$ due to sampling with $L_{\mathbf{b}}$. Equivalently, the corresponding component $x^{\mathbf{b}}(\mathbf{t})$ of $x(\mathbf{t})$ can be recovered by filtering the sampled image $x_{\mathbf{b}}(\mathbf{t})$ with an ideal bandpass filter with frequency support $\mathcal{A}^{\mathbf{b}}$.

Lemma 2: Suppose $M(B)$ is a Manhattan set and $x(\mathbf{t})$ is an image whose spectrum $X(\mathbf{u})$ is bandlimited to $\mathcal{M}(B)$. Then if $\mathbf{b}$ has maximal weight in $\overline{B}$,

$$X^{\mathbf{b}}(\mathbf{u}) \;=\; H^{\mathbf{b}}(\mathbf{u})\, X_{\mathbf{b}}(\mathbf{u})\,, \tag{17}$$

where $H^{\mathbf{b}}(\mathbf{u}) = 1$ for $\mathbf{u} \in A^{\mathbf{b}}$, and 0 otherwise.

Proof: Consider any $\mathbf{b}$ of maximal weight. Since, according to (4), $X_{\mathbf{b}}(\mathbf{u})$ consists of replicas of $X(\mathbf{u})$ at the frequencies in $L_{\mathbf{b}}^*$, since $X(\mathbf{u})$ can be decomposed into its components on Nyquist atoms $\{\mathcal{A}^{\mathbf{b}'} : \mathbf{b}' \in \{0,1\}^d\}$, and since $X(\mathbf{u})$ is bandlimited to $\mathcal{M}(B) = \bigcup_{\mathbf{b}' \in \overline{B}} \mathcal{A}^{\mathbf{b}'}$, it suffices to argue that for all $\mathbf{b}' \in \overline{B}$, no replica of $\mathcal{A}^{\mathbf{b}'}$ intersects $\mathcal{A}^{\mathbf{b}}$. First, Lemma 1(a) applied with $\mathbf{s} = \mathbf{b}' = \mathbf{b}$ shows that no replica of $\mathcal{A}^{\mathbf{b}}$ with $\mathbf{v} \neq \mathbf{0}$ intersects $\mathcal{A}^{\mathbf{b}}$. Second, Corollary 1 and the fact that $\mathbf{b}$ has maximal weight in $\overline{B}$ imply that for any other $\mathbf{b}' \in \overline{B}$, no replica of $\mathcal{A}^{\mathbf{b}'}$ can overlap $\mathcal{A}^{\mathbf{b}}$. □

Once $X(\mathbf{u})$ has been recovered in all such highest weight (highest frequency) M-atoms, the next lemma shows that $X(\mathbf{u})$ can then be recovered in the next highest weight M-atoms by canceling the contributions due to atoms with larger weight. In effect, the aliasing of one atom comes only from atoms with larger weight, i.e., higher frequency. Specifically, for any such $\mathbf{b}$, it shows that $X^{\mathbf{b}}(\mathbf{u})$ can be recovered from the spectrum $X_{\mathbf{b}}(\mathbf{u})$ due to sampling with $L_{\mathbf{b}}$ simply by first subtracting each replica $X^{\mathbf{b}'}(\mathbf{u} - \mathbf{v})$, $\mathbf{v} \in L_{\mathbf{b}}^*$, of every M-atom $\mathbf{b}'$ with $\|\mathbf{b}'\| > \|\mathbf{b}\|$, and then extracting the $\mathcal{A}^{\mathbf{b}}$ portion of the resulting "de-aliased" spectrum.

Lemma 3: Suppose $M(B)$ is a Manhattan set and $x(\mathbf{t})$ is an image whose spectrum $X(\mathbf{u})$ is bandlimited to $\mathcal{M}(B)$. If $X^{\mathbf{b}'}(\mathbf{u})$ is known for all $\mathbf{b}'$ larger than $\mathbf{b}$, then

$$X^{\mathbf{b}}(\mathbf{u}) = H^{\mathbf{b}}(\mathbf{u})\Big[X_{\mathbf{b}}(\mathbf{u}) \;-\; \sum_{\mathbf{b}':\, \|\mathbf{b}'\| > \|\mathbf{b}\|} X_{\mathbf{b}}^{\mathbf{b}'}(\mathbf{u})\Big]\,, \tag{18}$$

where $H^{\mathbf{b}}(\mathbf{u})$ is defined in the previous lemma,

Proof:

$$\begin{aligned}\text{RHS of (18)} &= H^{\mathbf{b}}(\mathbf{u})\Big[\sum_{\mathbf{v}\in L^*_{\mathbf{b}}} X(\mathbf{u}-\mathbf{v}) - \sum_{\mathbf{b}':\,\|\mathbf{b}'\|>\|\mathbf{b}\|}\ \sum_{\mathbf{v}\in L^*_{\mathbf{b}}} X^{\mathbf{b}'}(\mathbf{u}-\mathbf{v})\Big] \\ &= H^{\mathbf{b}}(\mathbf{u})\Big[\sum_{\mathbf{b}':\,\|\mathbf{b}'\|\le\|\mathbf{b}\|}\ \sum_{\mathbf{v}\in L^*_{\mathbf{b}}} X^{\mathbf{b}'}(\mathbf{u}-\mathbf{v})\Big] \\ &= X^{\mathbf{b}}(\mathbf{u})\,,\end{aligned} \tag{19}$$

where the first equality uses (4), the second uses (14), and the last derives from Lemma 1 and its corollary. In particular, for the $\mathbf{b}'=\mathbf{b}$ term in the above sum, Part (a) of Lemma 1 applied with $\mathbf{s}=\mathbf{b}'=\mathbf{b}$ shows that all replicas of $\mathcal{A}^{\mathbf{b}}$ (with $\mathbf{v}\neq\mathbf{0}$) do not overlap $\mathcal{A}^{\mathbf{b}}$ and, consequently, are eliminated by the filter $H^{\mathbf{b}}(\mathbf{u})$. Corollary 1 implies every replica of $\mathcal{A}^{\mathbf{b}'}$ (with $\mathbf{v}\neq\mathbf{0}$) does not overlap $\mathcal{A}^{\mathbf{b}}$ and, consequently, is again eliminated by the filter. Since also $\mathcal{A}^{\mathbf{b}'}$ does not overlap $\mathcal{A}^{\mathbf{b}}$, the only term in the sum not eliminated by the filter is $X^{\mathbf{b}}(\mathbf{u})$, which establishes (18). □

Note that the sum in (18) can be limited to $\mathbf{b}'\in\overline{B}$. Note also that Lemma 3 implies Lemma 2, because when $\mathbf{b}$ is a largest weight bi-step vector in $\overline{B}$, the summation term in (18) is zero, and so (18) reduces to (17).

An alternative way to write (18) is

$$X^{\mathbf{b}}(\mathbf{u}) = H^{\mathbf{b}}(\mathbf{u})\Big[X_{\mathbf{b}}(\mathbf{u}) - \sum_{\mathbf{b}':\,\|\mathbf{b}'\|>\|\mathbf{b}\|}\ \sum_{\mathbf{n}\in C_{\mathbf{b}}} X^{\mathbf{b}'}(\mathbf{u}-\mathbf{n}\odot\boldsymbol{\beta}_{\mathbf{b}})\Big] \tag{20}$$

where $\boldsymbol{\beta}_{\mathbf{b}}=(\beta_{\mathbf{b},1},\ldots,\beta_{\mathbf{b},d})$, with

$$\beta_{\mathbf{b},i} \;\triangleq\; \begin{cases}\frac{1}{\lambda_i}, & \text{if } b_i=1\\ \frac{1}{k_i\lambda_i}, & \text{if } b_i=0\end{cases}$$

and

$$C_{\mathbf{b}} \;\triangleq\; \big\{\mathbf{n}\in\mathbb{Z}^d : n_i=0 \text{ for } i \text{ s.t. } b_i=1, \text{ and} |n_i|\le k_i-1 \text{ for } i \text{ s.t. } b_i=0\big\}.$$

To demonstrate (20), we note that since all atoms of the Manhattan partition are contained in $\mathcal{N}_{\mathbf{1}}$, one can eliminate from the last sum in (19) any $\mathbf{v}$ such that $(\mathcal{N}_{\mathbf{1}}+\mathbf{v})\cap\mathcal{N}_{\mathbf{1}}=\emptyset$. This leads to limiting the sum to $\mathbf{v}$ such that $|v_i|<\frac{1}{\lambda_i}$ for each i. Taking into account what $\mathbf{v}$'s are in $L^*_{\mathbf{b}}$ leads to (20). For the usual 2D case, in which $B=\{(1,0),(0,1)\}$, (20) gives a different reconstruction formula than in Section III for the spectrum in the coarse Nyquist region, $X^C(\mathbf{u})=X^{(0,0)}(\mathbf{u})$ (see (5)-(7)). Specifically, it subtracts terms involving both $X^V(\mathbf{u})=X^{(0,1)}(\mathbf{u})$ and $X^H(\mathbf{u})=X^{(1,0)}(\mathbf{u})$ from $X_C(\mathbf{u})$, whereas the formula in Section III subtracts terms involving $X^H(\mathbf{u})$ from $X_V(\mathbf{u})$. Moreover, the summation over $\mathbf{n}$

in (20) sums over approximately twice as many values of $\mathbf{v}$. This is because it conservatively includes all $\mathbf{v} \in L_{\mathbf{b}}^*$ such that $\mathcal{N}_1 + \mathbf{v} \cap \mathcal{N}_1 \neq \emptyset$, whereas the formula in Section III includes only $\mathbf{v}$'s such that $\mathcal{N}_V + \mathbf{v} \cap \mathcal{N}_C \neq \emptyset$. If desired $C_{\mathbf{b}}$, in (20) could be replaced by a smaller set $C_{\mathbf{b},\mathbf{b}'}$ that depends on $\mathbf{b}'$ as well as $\mathbf{b}$.

The basic idea behind following theorem, which is the main result of this section, is that the process of finding $X^{\mathbf{b}}(\mathbf{u})$ for smaller and smaller weight $\mathbf{b}$'s can continue until $X^{\mathbf{0}}$, the spectrum in $\mathcal{A}^{\mathbf{0}}(\mathbf{u}) = \mathcal{N}_{\mathbf{0}}$, is found, and all of $X(\mathbf{u})$ is known. As a result, $x(\mathbf{t})$ will also be known.

Theorem 3: Multidimensional Manhattan Sampling Theorem. Suppose we sample an image $x(\mathbf{t})$ with Manhattan set $M(B)$. If the image spectrum $X(\mathbf{u})$ is bandlimited to $\mathcal{M}(B)$, then for each $\mathbf{b} \in \overline{B}$, $X^{\mathbf{b}}(\mathbf{u})$ can be exactly recovered from the samples with the following "onion-peeling" approach — apply Lemma 2 for the largest $\mathbf{b}$'s in $\overline{B}$, and then repeatedly apply Lemma 3 for the next largest $\mathbf{b}$'s. Then, $x(\mathbf{t})$ can be exactly recovered from

$$x(\mathbf{t}) \;=\; \mathcal{F}^{-1}\Big\{ \sum_{\mathbf{b}\in\overline{B}} X^{\mathbf{b}}(\mathbf{u}) \Big\}\,.$$

Proof: From (14) and the bandlimitation of $X(\mathbf{u})$, it is clear that $X(\mathbf{u})$ can be recovered if $X^{\mathbf{b}}(\mathbf{u})$ is recovered for each $\mathbf{b} \in \overline{B}$. First, $X^{\mathbf{b}}(\mathbf{u})$ can be recovered via Lemma 2 for the largest $\mathbf{b}$'s in $\overline{B}$, which correspond to the highest frequency M-atoms. Next, repeatedly applying Lemma 3 enables one to recover $X^{\mathbf{b}}(\mathbf{u})$ for the Nyquist atoms corresponding to the largest of the remaining $\mathbf{b}$'s, until $X^{\mathbf{0}}(\mathbf{u})$, corresponding to the lowpass atom, is recovered. □

While this theorem indicates a frequency domain reconstruction, followed by an inverse transform, a direct spatial domain reconstruction is also possible. As we now delineate, this involves reconstructing each $x^{\mathbf{b}}(\mathbf{t})$, $\mathbf{b} \in \overline{B}$, and then using

$$x(\mathbf{t}) \;=\; \sum_{\mathbf{b}\in\overline{B}} x^{\mathbf{b}}(\mathbf{t})\,.$$

Taking the inverse transform of (20) yields

$$\begin{aligned}
x^{\mathbf{b}}(\mathbf{t}) \;&=\; h^{\mathbf{b}}(\mathbf{t}) \star \Big[x_{\mathbf{b}}(\mathbf{t}) \;-\; \sum_{\|\mathbf{b}'\|>\|\mathbf{b}\|} x_{\mathbf{b}}^{\mathbf{b}'}(\mathbf{t}) \Big] \\
&=\; h^{\mathbf{b}}(\mathbf{t}) \star \Big[K(\mathbf{b}) \sum_{\mathbf{t}'\in L_{\mathbf{b}}} \delta(\mathbf{t}-\mathbf{t}') \Big(x(\mathbf{t}) - \sum_{\|\mathbf{b}'\|>\|\mathbf{b}\|} x^{\mathbf{b}'}(\mathbf{t}) \Big) \Big] \\
&=\; K_{\mathbf{b}} \sum_{\mathbf{t}'\in L_{\mathbf{b}}} \Big(x(\mathbf{t}') \;-\; \sum_{\|\mathbf{b}'\|>\|\mathbf{b}\|} x^{\mathbf{b}'}(\mathbf{t}') \Big)\, h^{\mathbf{b}}(\mathbf{t}-\mathbf{t}')\,,
\end{aligned}$$

where $h^{\mathbf{b}}(\mathbf{t}) = \mathcal{F}^{-1}\{H^{\mathbf{b}}(\mathbf{u})\}$, $\star$ denotes convolution, and $K_{\mathbf{b}} = \prod_{i:b_i=1} \lambda_i \times \prod_{i:b_i=0} k_i\lambda_i$. This shows how $x^{\mathbf{b}}(\mathbf{t})$ can be found — first for the largest $\mathbf{b}$'s from samples of $x(\mathbf{t})$ taken on $L_{\mathbf{b}}$, then for the next largest $\mathbf{b}$'s from samples of $x(\mathbf{t})$ taken on $L_{\mathbf{b}}$, as well as samples of $x^{\mathbf{b}'}(\mathbf{t})$ taken on $L_{\mathbf{b}}$ for all larger $\mathbf{b}'$, and so on. It remains to find a formula for $h^{\mathbf{b}}(\mathbf{t})$.

To find a formula for $h^{\mathbf{b}}(\mathbf{t})$, which is the inverse transform of $H^{\mathbf{b}}(\mathbf{u})$, which in turn has support $\mathcal{A}^{\mathbf{b}}$, we begin by recalling that $\mathcal{A}^{\mathbf{b}}$ is the union of $2^{\|\mathbf{b}\|}$ orthotopes in frequency space. Along each dimension i, these orthotopes are centered at zero if $b_i = 0$, and at $\pm c_i$ if $b_i = 1$, where

$$c_i \triangleq \frac{1}{2}\Big(\frac{1}{2\lambda_i} + \frac{1}{2k_i\lambda_i}\Big).$$

Additionally, along the ith dimension, these orthotopes have length

$$w_i(\mathbf{b}) \triangleq \begin{cases} \frac{1}{2\lambda_i} - \frac{1}{2k_i\lambda_i}, & b_i = 1 \\ \frac{1}{k_i\lambda_i}, & b_i = 0. \end{cases}.$$

Using these quantities, we can write the filter $H^{\mathbf{b}}(\mathbf{u})$ as

$$H^{\mathbf{b}}(\mathbf{u}) = \Big[\prod_{i=1}^{d} \mathrm{rect}\Big(\frac{u_i}{w_i(\mathbf{b})}\Big)\Big] \star \Big[\prod_{i:b_i=0} \delta(u_i) \prod_{i:b_i=1} \Big[\delta(u_i - c_i) + \delta(u_i + c_i)\Big]\Big],$$

where $\mathrm{rect}(x) \triangleq 1$ for $|x| < \frac{1}{2}$, and 0 otherwise. Note that the first term is an orthotope centered at $\mathbf{0}$, and the convolution with delta functions shifts the orthotopes along all dimensions i such that $b_i = 1$. Taking the inverse transform yields

$$h_{\mathbf{b}}(\mathbf{t}) = \prod_{i=1}^{d} w_i(\mathbf{b}) \,\mathrm{sinc}(w_i(\mathbf{b})t_i) \cdot \prod_{i:b_i=1} 2\cos(2\pi c_i t_i),$$

where $\mathrm{sinc}(t) \triangleq \frac{\sin \pi t}{\pi t}$. Observe that, as mentioned in the introduction, these impulse responses depend on the k_i's and λ_i's, but not the choice of bi-step lattices that comprise the Manhattan set. Moreover, the λ_i's have only a simple spatial scaling effect on the filters.

I. Achievement of Landau lower bound on sampling density

We now show that the volume of $\mathcal{M}(B)$, denoted $|\mathcal{M}(B)|$, equals the sampling density of $M(B)$. As a result, the set of images bandlimited to $\mathcal{M}(B)$ is a maximal set of images that are reconstructable from sample set $M(B)$. Equivalently, $M(B)$ has the smallest density of any sampling set such that all images bandlimited to $\mathcal{M}(B)$ are reconstructable.

Since the M-atoms partition $\mathcal{M}(B)$, we can calculate $|\mathcal{M}(B)|$ simply by summing over the volumes of the M-atoms $\mathcal{A}^{\mathbf{b}}$ for $\mathbf{b} \in \overline{B}$:

$$|\mathcal{M}(B)| = \sum_{\mathbf{b}\in\overline{B}} |\mathcal{A}^{\mathbf{b}}| = \sum_{\mathbf{b}\in\overline{B}} \Big(\prod_{i:b_i=1} 2\Big(\frac{1}{2\lambda_i} - \frac{1}{2k_i\lambda_i}\Big) \prod_{i:b_i=0} \frac{1}{k_i\lambda_i} \Big)$$
$$= \frac{\sum_{\mathbf{b}\in\overline{B}} \prod_{i:b_i=1}(k_i-1)}{\prod_{i=1}^{d} k_i\lambda_i}$$

Comparing the above to (12), we see that $|\mathcal{M}(B)|$ equals the sampling density $\rho(B)$.

J. Discrete-space images

The d-dimensional Manhattan sampling theorem and reconstruction procedures can be straightforwardly extended to discrete-space images in d dimensions in the same fashion as for two dimensions. For example, for infinite-support images, frequencies need to be scaled by 2π, and for finite-support images, each spatial resolution T_i must be a multiple of $k_i\lambda_i$ and the Manhattan atoms $\widetilde{A}_{\mathbf{b}}$ need to be redefined to be consistent with the DFT, as was done for the discrete Nyquist region $\widetilde{\mathcal{N}}_{\alpha_1,\alpha_2}$. Here, we simply give the main step of the frequency-space onion-peeling reconstruction algorithm for reconstructing a Manhattan bandlimited discrete-space image $x[\mathbf{t}]$ with finite support sampled with Manhattan set $M(B)$:

$$X^{\mathbf{b}}[\mathbf{t}] \;=\; \widetilde{H}^{\mathbf{b}}[\mathbf{u}] \;\mathrm{DFT}\Big\{ x_{\mathbf{b}}[\mathbf{t}] \;-\; \sum_{\mathbf{b}':\|\mathbf{b}'\|>\|\mathbf{b}\|} x_{\mathbf{b}}^{\mathbf{b}'}[\mathbf{t}] \Big\},$$

where $x_b[\mathbf{t}]$ and $x_{\mathbf{b}}^{\mathbf{b}'}[\mathbf{t}]$ denote, respectively, the $L_{\mathbf{b}}$ subsamplings of the Manhattan samples (scaled by $K_{\mathbf{b}}$), and the previously reconstructed atom $x^{\mathbf{b}'}[\mathbf{t}]$, and $\widetilde{H}^{\mathbf{b}}[\mathbf{u}]$ denotes an ideal bandpass filter for atom $\widetilde{\mathcal{A}}_{\mathbf{b}}$.

V. Concluding Remarks

In two dimensions, this paper has shown that from samples of a Manhattan set one can perfectly reconstruct any image that is bandlimited to the union of the Nyquist regions of the horizontal and vertical rectangular lattices comprising the Manhattan set. It also prescribed a straightforward linear reconstruction procedure, for continuous- and discrete-space images.

For three and higher dimensions, this paper has identified Manhattan sets as the union of a finite number of bi-step rectangular lattices, with the result that many Manhattan geometries are possible. It introduced an efficient binary-vector representation of bi-step lattices, and consequently Manhattan sets, which enabled the specification of a partition of the dense rectangular lattice into collections of cosets of the coarse lattice. This, in turn, enabled the density of a Manhattan to be computed. The representation

of bi-step lattices also enabled a partition of the Nyquist region of the dense rectangular lattice, which in turn enabled a precise analysis of the aliasing, i.e., spectral overlaps, by the atoms of any particular type in the spectral replicas induced by any particular bi-step lattice subsampling. With this, it was shown that images bandlimited to the union of the Nyquist regions of the bi-step lattices comprising the Manhattan set can be perfectly reconstructed using an an efficient closed-form onion-peeling type reconstruction algorithm that reconstructs the image spectrum working from higher to lower frequency atoms of the partition. At each step, the algorithm works with samples of one particular bi-step lattice (of the Manhattan set), and obtains the spectrum of the image in the corresponding atom of the frequency partition by subtracting contributions due to aliasing of previously determined atoms of the partition. Both frequency- and time-domain versions of the algorithm were given. It was also shown that the set of Manhattan bandlimited images is maximal in the Landau sense. To the best of our knowledge, this is the first demonstration that images bandlimited to the union of Nyquist regions can be recovered from the union of the corresponding lattices.

There are several avenues for future research. One could seek to extend the results to continuous-space images whose spectra contain delta functions, e.g. periodic images. Second, instead of a recursive onion-peeling reconstruction, one could seek direct closed-form linear reconstructions, as in [20], which might be useful for implementations, though they might have less intuitive appeal. This is not difficult in 2D, but is more challenging in higher dimensions. Finally, whereas M-sampling can be viewed as sampling (in various ways) on the boundaries of a rectangular (hyper-rectangular) lattice tessellation, one could seek sampling theorems and reconstruction procedures for images sampled on the boundaries of other lattice tessellations, such as a hexagonal tessellation.

REFERENCES

[1] M.G. Reyes, X. Zhao, D.L. Neuhoff, T.N. Pappas, "Lossy compression of bilevel images based on Markov random fields," *Proc. IEEE ICIP*, San Antonio, pp. II-373–376, Sept. 2007.

[2] M.G. Reyes and D.L. Neuhoff, "Arithmetic encoding of Markov random fields," *Proc. IEEE Int. Symp. Inf. Thy.*, Seoul, pp. 532–536, June 2009.

[3] M.G. Reyes and D.L. Neuhoff, "Lossless reduced-cutset coding of Markov random fields," *Proc. Data Compression Conf.*, Snowbird, pp. 386–395, Mar. 2010.

[4] A. Farmer, A. Josan, M.A. Prelee, D.L. Neuhoff, and T.N. Pappas, "Cutset sampling and reconstruction of images," *Proc. IEEE ICIP*, Brussels, pp. 1949–1952, Sept. 2011.

[5] M.A. Prelee, D.L. Neuhoff, and T.N. Pappas, "Image reconstruction from a Manhattan grid via piecewise plane fitting and Gaussian Markov random fields," *Proc. IEEE ICIP*, Orlando, FL, pp. 2061–2064, Sept. 2012.

[6] M.A. Prelee and D.L. Neuhoff, "Image interpolation from Manhattan cutest samples via orthogonal gradient method, "*Proc. IEEE ICIP*, Paris, Fr, pp. 1842–1846, Oct. 2014.

[7] M.A. Prelee and D.L. Neuhoff, "Energy efficient source localization on a manhattan grid wireless sensor network," in *Proc. IEEE ICASSP*, pp. 4266–4270, Vancouver, May 2013.

[8] M.A. Prelee and D.L. Neuhoff, "Performance–energy tradeoffs in cutset wireless sensor networks," to be presented at *IEEE ICASSP*, May 2014.

[9] J. Unnikrishnan and M. Vetterli, "Sampling high-dimensional bandlimited fields on low-dimensional manifolds," *IEEE Trans. Info. Thy.*, vol. 59, pp. 2103–2127, April 2013.

[10] J. Unnikrishnan and M. Vetterli, "Sampling and reconstruction of spatial fields using mobile sensors," *IEEE Trans. Sig. Proc.*, vol. 61,pp. 2328–2340, May, 2013.

[11] D.P. Peterson and D. Middleton, "Sampling and reconstruction of wave-number limited functions in N-dimensional Euclidean spaces," *Inf. Control*, vol. 5, no. 4, pp. 279–323, 1962.

[12] A. Rosenfeld and A.C. Kak, *Digital Picture Processing*. New York, NY: Academic Proess, 1976.

[13] A.M. Tekalp, *Digital Video Processing*. Upper Saddle River, NJ: Prentice-Hall, 1995.

[14] D.E. Dudgeon and R.M. Meresereau, *Multidimensional Digital Signal Processing*. Englewood Cliffs, N.J.: Prentice-Hall, 1984.

[15] H. Landau, "Necessary density conditions for sampling and interpolation of certain entire functions," *Acta Math*, vol. 117, no. 1, pp. 37–52, 1967.

[16] N.T. Gaarder, "A note on the multidimensional sampling theorem," *Proc. IEEE*, vol. 60, pp. 247–248, Feb. 1972.

[17] R. J. Marks II, "Multidimensional-signal sample dependency at Nyquist densites," *J. Opt. Soc. Am. A*, Vol. 3, pp. 268–273, Feb. 1986.

[18] K.F. Cheung and R.J. Marks II, "Image sampling below the Nyquist density without aliasing," *J. Opt. Soc. Amer.*, vol. 7, pp. 92–105, Jan. 1990.

[19] K.F. Cheung, "A multidimensional extension of Papoulis' generalized sampling expansion with the application in minimum density sampling," *Advanced Topics in Shannon Sampling and Interpolation Theory*. New York: Springer-Verlag, pp. 85–119, 1993.

[20] A. Faridani, "An application of a multidimensional sampling theorem to computed tomography," *Contemp. Math.*, vol. 113, pp. 65-80, June 1990.

[21] A. Faridani, "A generalized sampling theorem for locally compact Abelian groups," *Math. Comp.*, vol. 63, no. 207, pp. 307–327, 1994.

[22] D. Walnut,"Nonperiodic sampling of bandlimited functions on union of rectangular lattices," *J. Four. Anal. Appl.*, vol. 2, no. 5, pp. 435-452, 1996.

[23] R. Venkataramani and Y. Bresler, "Perfect reconstruction formulas and bounds on aliasing error in sub-Nyquist nonuniform sampling of multiband signals," *IEEE Trans. Inf. Thy.*, vol. 46, pp. 2173–2183, Sept. 2000.

[24] R. Venkataramani and Y. Bresler, "Optimal sub-Nyquist nonuniform sampling and reconstruction for multiband signals," *IEEE Trans. Sig. Proc.*, vol. 49, pp. 2301–2313, Oct. 2001.

[25] H. Behmard and A. Faridani, "Sampling of bandlimited functions on unions of shifted lattices," *J. Four. Anal. Appl.*, v. 8, 1, pp. 43–58, 2002.

[26] H. Behmard, "Efficient Reconstruction Algorithms Using Shifted Lattices," *IEEE Trans. Sig. Proc.*, vol. 57, pp. 2548–2557,

July 2009.

[27] J. Unnikrishnan and M. Vetterli, "On sampling a high-dimensional bandlimited field on a union of shifted lattices," *in Proc. IEEE Int. Symp. Inf. Thy.* pp. 1473-1477, Cambridge, MA, July 2012.

[28] A. Papoulis, "Generalized sampling expansion," *IEEE Trans. Circuits and Systems*, vol. 24, pp. 652–654, Nov. 1977.

[29] N.P. Willis and Y. Bresler, "Optimal scan design for time varying tomographic imaging I: Theoretical analysis and fundamental limitations," *IEEE Trans. Image Proc.*, vol. 4, pp. 642–654, May 1995.

[30] "Multilattice sampling strategies for region of interest dynamic MRI," G. Rilling, Y. Tao, I. Marshall, and M.E. Davies. *Magnetic Resonance in Medicine*, vol. 70, no. 2, pp. 392–403, Aug. 2013.

[31] J.R. Higgins, "Five short stories about the cardinal series," *Bull. AMS*, vol. 12, no. 1, pp. 45–89, 1985